Study of non-neutral electron plasma lenses for focusing laser-driven ion beams

Titus-Stefan Dascalu

Department of Physics
Imperial College London

*Thesis submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy at Imperial College London*

May 2023

Abstract

As the need for scale-up in the provision of radiotherapy (RT) is clear, developing cost-effective and robust technologies for RT is ever more crucial. Novel proton- and ion-acceleration techniques based on high-power lasers can provide beams with unique characteristics. To make full use of a number of these techniques, the particles need to be converted efficiently into a beam.

This thesis describes work on the development of space-charge plasma lenses for capturing and focusing laser-driven protons and ions. A non-neutral electron plasma, which is confined by external electric and magnetic fields, generates transverse focusing. The work is undertaken as part of the LhARA project.

As a starting point, an existing prototype Gabor lens is modelled with particle-in-cell (PIC) simulations. It is found that an unstable plasma can lead to deleterious focusing effects comparable to those recorded previously in beam tests. A number of plasma instabilities is linked to the initial production of non-uniform electron clouds within the lens. Validation of the PIC code used to model confined plasmas is carried out using the results from experiments with low-density, magnetically-trapped electron clouds. The capabilities of the particle trap used throughout the experiments are characterised by measurements of the density, size, and evolution in time of the confined electron clouds.

This thesis also includes a first design of a normal-conducting solenoid suitable for LhARA. The magnet requires moderate to large power consumption due to Joule losses, in addition to a complex cooling system. These requirements are found to be particularly sensitive to the bore diameter and a strategy to reduce them with a transversely-graded coil is described.

Finally, the performance of plasma lenses and solenoids is evaluated by tracking a realistic laser-driven proton beam through the LhARA Stage 1 beam-line. Two models of the space-charge field produced by plasma lenses are presented.

Acknowledgements

First, I would like to express my immense gratitude to my supervisor, Ken Long, who guided me with unwavering support during my time both at Imperial and away from the college. You always granted me the independence to pursue my own ideas and interests, for which I am extremely grateful. It has been a source of inspiration to constantly discover how your thoughtful, attentive, and deeply-caring spirit can turn so many wheels of scientific collaboration.

I would also like to extend my sincere thanks to Chris Baker who offered a great amount of assistance and hands-on help during my work in Swansea. The practical suggestions and insightful criticism you extended to me taught me about the thoroughness of experimental investigation, while your patience cannot be underestimated. Another very special word of thanks goes to Elisabetta Boella for the data she provided and, more importantly, for her advice and the scientific discussion we had.

For sharing their extensive knowledge and helpful guidance at several turning points of my work, I would like to thank Colin Whyte, Mike Charlton, Bob Bingham, and William Shields. Thanks must go to the Imperial Accelerator group, especially Jaroslaw Pasternak and Ajit Kurup, for their valuable collaboration. A huge thank you to Maria Maxouti and Josie McGariggle for making the 11th floor office in Blackett even more enjoyable, and to Rebecca Taylor for sharing her excitement in accelerator physics during my stay at CERN. Moreover, I gratefully acknowledge the funding support that I received trough the Class of 1964 scholarship, Imperial College London.

I am especially indebted to my parents for their continuous and unconditional support, not only during my thesis but already during my studies and life in general. On the same note, thank you to my brother for reminding me to find joy in all aspects of life.

Above all, I am deeply grateful to Raluca Dascalu. You have showed boundless personal strength and independence while you have never stopped motivating and supporting me. I cannot thank you enough.

Statement of Originality

I hereby declare that the material presented in this thesis is my own work unless otherwise stated or referenced. Some figures referenced as from other sources were created by me, but appear in other public documents.

The analysis and simulations described in chapter 4 were my own work. The prototype of the Gabor lens at Imperial College London was built and tested prior to the start of my PhD programme. I produced all the simulations and the models for comparison to data, while the data and the beam images were collected in a beam test in which I was not involved.

The measurements described in chapter 5 and the data that are presented are the result of my own work. Changes to the experimental apparatus prior to different sets of measurements were done either by myself alone or under the supervision of Dr. Chris Baker from Swansea University who contributed entirely on his own to the modification of the electron source, its installation onto the beam-line, and maintaining the proper vacuum conditions along the beam-line during the measurements.

The design study described in chapter 6 was entirely my own work and I was solely responsible for producing all the figures shown in this chapter. The new models of the plasma lens outlined in chapter 7 were created, implemented, and tested by myself. The beam-tracking simulations described in chapter 7 were run and analysed by myself and were based on models of the LhARA beam-lines to which multiple members of the LhARA Collaboration have contributed.

All of the work outlined above was continuously reported to the LhARA Collaboration and the CCAP group at Imperial College London.

As well as the work described in this thesis, I also worked for six months at CERN as part of the BE-ABP-HSL section. I investigated preliminary design ideas for an electrostatic low energy beam transport line (ELEBT) for use in future high-current H^- beam injectors. The corresponding work

is not presented in this thesis because it is not thematically linked to the development of electron plasma lenses for LhARA.

Titus-Stefan Dascalu

May 2023

Copyright Statement

Contents

List of Figures

List of Tables

Introduction

Cancer is the second most common cause of death globally [1], causing nearly one in six deaths. The incidence of new cases is predicted to increase globally due to population growth and ageing from 17 million new cases per year at present to 28 million new cases per year by 2040 [2]. Many cancers can be cured if detected early and treated effectively, and there are many different approaches to treating cancer. Amongst these approaches, radiotherapy (RT) is used in treatments of more than 50% of cancer patients and results in 40% of cancer cures, second only to surgery [3]. In Western countries, even a larger fraction of cancer patients (approximately two thirds) receive radiotherapy at some stage of their treatment [4]. It is estimated that more than two million people will require RT in Europe alone in 2025 [5] and that 27 million life-years could be save in low- and middle-income countries if the access to RT could be scaled up [6].

Radiation therapy is an advancing treatment modality of which outcomes have seen substantial improvements in recent years. For example, tumour control can reach up to 80% in some cases of head and neck cancers [4]. Nevertheless, many cancer types still respond poorly to radiotherapy treatments as is the case of hypoxic and other radioresistant tumours. The efficacy of RT is mainly limited by the normal tissue tolerance and radiation-induced adverse effects [4]. In particular, paediatric patients are known to be severely affected by long-term consequences of RT which impacts numerous functional domains and quality of life [7]. Research in RT is therefore focused on improving the current treatment strategies or finding new ones to increase the rates of tumour control while simultaneously improving the sparing of normal tissue.

An important role in this context has been played by proton and ion beams. While RT is most frequently delivered with the use of photons (X-rays), there is an increasing emphasis on the use and further development of proton- and ion-beam therapy. Compared to photons, proton and ion beams exhibit a more localised energy deposition which allows a more conformal irradiation of the target and sparing of healthy tissue. Moreover, a lower radiation dose delivered by heavy ions is needed to produce the same therapeutic effect that would be obtained using photons. This makes heavy ions suitable candidates for the treatment of hypoxic and other radioresistant tumours [8]. However, there are still significant biological uncertainties in the impact of ion beams on living tissue.

Systematic studies are needed to improve the understanding of how the physical characteristics of particle beams generate the biological response of healthy and tumour tissues to ionising radiation. The Laser-hybrid Accelerator for Radiobiological Applications (LhARA) is a planned particle beam facility [9, 10] dedicated to such systematic radiobiological research. The LhARA facility has been conceived as a new, highly flexible source

of radiation that can contribute to studies of radiobiology with the exploitation of ultra-high dose rates, multiple ion species, and novel temporal- and spatial-fractionation schemes. The vision of the LhARA collaboration is to contribute to the development of better patient-specific treatment plans that correctly calculate the dose required to treat a particular tumour while minimising the damage to the healthy tissue and organs at risk. Furthermore, LhARA has been designed to allow the study of the potential therapeutic benefit of beams with properties that are not available at conventional facilities.

The beam properties that could lead to an improvement in the sparing of healthy tissue are dictated by recent reports which provide exciting evidence of therapeutic benefit when radiation is delivered in new regimes. Particle beam therapy (PBT) today is delivered at low dose rates ($<10\,\mathrm{Gy/min}$) within restricted beam characteristics by employing a small number of temporal schemes and spatial distributions. The damage to the healthy tissue limits the dose delivered and, thus, constraints the clinical efficacy. Evidence of enhanced therapeutic benefit has been recently found with the use of novel beams with strikingly different characteristics, e.g. very high dose per fraction, very high dose rate ($>40\,\mathrm{Gy/s}$, "FLASH") [11–13], and spatially fractionated dose from "mini-beams" (MBRT) [14, 15]. The exploration of new regimens of PBT now provides the impetus for a radical transformation of PBT.

In addition to the programme of radiobiology, LhARA aims to prove accelerator technologies required to deliver the necessary scale-up in the provision of PBT. The Particle Therapy Co-Operative Group (PTCOG) [16] currently lists 105 proton- and 14 carbon-ion-therapy facilities worldwide [17]. As these facilities are located predominantly in high-income countries, nearly 70% of cancer patients in low- and middle-income countries globally do not have access to any modality of radiation therapy [3]. Furthermore, less than 1% of RT patients worldwide are treated with protons or other ions [18] despite their widely recognised clinical advantages over conventional X-rays. In this context, the main technologies that are proposed by the LhARA collaboration and that have the potential to decrease the cost and size of a clinical facility are: a laser-driven proton and ion source, electron-plasma lenses and compact RF post-acceleration.

Laser-driven proton and ion sources are an emerging technology and have been proposed as an alternative to conventional proton and carbon-ion facilities for RT [19–21] for a series of reasons. In addition to the relative compactness of laser systems, the beams generated by high power laser have unique properties [22] that allow the investigation of new regimes of RT. Radiobiology research with laser-driven protons and ions has been planned or conducted at a number of laboratories [23, 24] and, more recently, research platforms based on petawatt lasers have started to serve pilot studies on *in-vitro* cell cultures [25] and small-animal models [26, 27]. The most widely used acceleration mechanism based on high-power lasers is known as Target Normal Sheath Acceleration (TNSA). In recent years, the TNSA mechanism has been developed and significant progress has been made which has turned it into a robust and effective technique for use in laser-driven ion sources.

In LhARA, the laser-driven source is envisioned to generate multiple ion species forming on-demand beams with a variety of time, spatial, and spectral structures. Moreover, the high initial ion energies ($\gg 1\,\mathrm{MeV/u}$) help to overcome the space-charge limitations of conventional sources and together with the high flux of ions

produced at the source and efficient capture can lead to the delivery of radiation at ultra-high instantaneous dose rates ($\sim$70–700 Gy/s [9]). However, laser-driven ion sources based on the TNSA mechanism create beams that are highly divergent and have a large energy spread. These issues need to be addressed by effective strategies for capturing the particles at the source and reducing the beam energy spread without significant loss of particles. If such a strategy is proved to work at the repetition rates of modern high-power lasers, 1–10 Hz, full use could be made of the unique properties of laser-generated proton and ion sources.

Conventionally, charged particle beams are focused by electrostatic or magnetic fields and the choice of a particular focusing system is based on the required focusing strength which, in turn, depends on the energy or momentum of the beam in most cases. Laser-driven sources typically produce particles with kinetic energy above 1 MeV–a threshold around which magnetic focusing becomes more efficient over electrostatic focusing, resulting in shorter focal lengths. Thus, magnetic focusing has been explored in association with laser-driven protons through systems based on quadrupole magnets, both electromagnetic [28, 29] and permanent (PMQ) [30–32], pulsed high-field solenoids [26, 33–35], and active plasma lenses (APL) [25]. Indeed, all of these devices allowed protons to be captured, focused, and collimated with variable efficacy and experiments have been successfully performed in which protons from laser-driven sources were transmitted to targets several meters downstream. Quadrupole magnets focus in one transverse plane and defocus the beam in the other, typically resulting in long beam-lines. Furthermore, achieving large field gradients in quadrupole magnets require small apertures which limit the capture efficiency–an effect especially significant when using PMQs. Even though pulsed high-field solenoids were demonstrated to allow better transmission of the beam [26], the repetition rate is currently limited to approximately 0.1 Hz [36]. Lastly, active plasma lenses with radii larger than one centimetre were found to degrade the beam quality due to the z-pinch effect in the plasma and, in turn, produce non-linear focusing [37, 38]. More recently, compact APLs based on gas-filled capillaries with diameters up to 2 mm have been demonstrated to generate stable focusing magnetic fields [39]. Nonetheless, the limited transverse size and length of the plasma currently lead to comparatively low proton capture efficiencies [25].

A possible route to overcome the limitations of conventional electrostatic and magnetic focusing devices in the context of laser-driven proton and ion beams is to employ the space-charge force of a non-neutral plasma to generate the desired focusing effect. There have been numerous attempts [40–48] to create a practical space-charge (plasma) lens[1] based on the original design proposed by Gabor [49] and the majority of these were able to produce a plasma and observe a focused ion beam. However, most of the devices constructed were capable of producing weakly-confined plasmas in quasi-steady states. While such plasmas can produce suitably high electron densities to satisfy the requirements for a space-charge lens, the lack of control over the plasma production can lead to deleterious effects. Some of the challenges that were reported are: the generation of a highly homogeneous plasma [51], the onset of plasma instabilities [52, 53], and the lens causing beam emittance

---

[1]If not specified otherwise, a *plasma lens*, *space-charge lens*, or *Gabor lens* will refer in this work to a non-neutral plasma that generates a focusing force on a charged-particle beam which propagates through the cloud of charge represented by the plasma. It will be generally considered that the plasma is confined by external magnetic and electric fields in a device with similar configuration to the first proposal of the lens by D. Gabor [49] and more recently known as a Penning-Malmberg trap [50].

growth [43, 44, 46, 51].

The Gabor lenses are an important part of the LhARA facility concept. The large divergence angle and energy spread of the protons or ions at the source are addressed in the conceptual design of LhARA [9] through the use of plasma lenses to provide strong focusing and to allow energy selection. Furthermore, the lens is suitable for operation at high repetition rate, $1–10\,\text{Hz}$, and, in principle, the focusing strength can be tuned by changing the strength of the fields that confine the plasma. The development of a plasma lens for LhARA started at Imperial College London with the construction and testing of a prototype Gabor lens [54] that could generate a high-density electron plasma ($\sim 1 \times 10^{14}\,\text{m}^{-3}$). The prototype lens was designed to confine the plasma using an axial magnetic field ($\leq 60\,\text{mT}$) and a configuration of electrodes biased at high voltages ($\leq 20\,\text{kV}$).

Despite laboratory investigations of the plasma produced by the prototype and a couple of tests of the lens with proton beams, the operation of the lens had limited success. A detrimental focusing effect of the lens on proton beams had been repeatedly recorded and the underlying source had not been identified. The results of the beam tests motivated a study of the plasma behaviour in unstable regimes and showed that significant challenges related to the production and stability of the plasma still need to be overcome for the lens to be used reliably in a research or treatment facility. Furthermore, it was understood that a key factor for the development of a new, improved prototype lens is a simulation method that can predict the most important aspects of the behaviour of the plasma confined by the lens prior to the construction of the physical prototype. More recently, a plan for a series of experiments was included in the preliminary phase of the R&D programme for LhARA [10]. The experiments are dedicated to the study of non-neutral plasmas with long storage times and radii large compared to the confining apparatus.

The central subject of this thesis is, therefore, the investigation of non-neutral electron plasma lenses as a modality of capturing and focusing protons and ions generated by laser-driven sources.

In the following, a brief summary is given of the work performed during the PhD programme.

Summary

The main work performed during this PhD concerned the identification and study of the limitations and the capabilities of non-neutral electron plasma (Gabor) lenses as devices to capture and focus laser-driven proton and ion beams. As the starting point of these investigations, the source of the deleterious focusing effect recorded in the beam tests of the first lens prototype built at Imperial College London was considered. For this, a model of the prototype lens was created and the corresponding electron plasma simulated with the particle-in-cell (PIC) method and a dedicated software. The results showed that an unstable plasma with the bulk of the electrons displaced from the central axis of the lens and rotating around it can convert thin proton pencil-beams into rings. A simplified analytical model was created for a plasma column that rotates around the central axis of the lens, known as the *diocotron mode*, and particle-tracking simulations showed that such a plasma leads to the formation of rings with similar size, aspect, and intensity distribution as those recorded in a beam test of the prototype lens. In order to reproduce qualitatively the ring patterns, the plasma density and period of rotation were tuned and the resulting range for the electron density represents an evaluation of the capability of the prototype lens to produce relatively high-density plasmas. The onset of an unstable plasma and its rotation about the axis of the lens, and, in turn, the deleterious focusing effect were attributed to the unsophisticated electron-production mechanism and to the lack of control over the initial distribution of the plasma inside the lens.

Based on the confirmed capability of particle-in-cell simulations to make qualitative predictions on the stability and evolution of the plasma contained by a Gabor lens, a series of experiments were planned to validate the PIC method and the particular software used in this work against measurements. A beam-line at Swansea Unviersity consisting of an electron source, a plasma trap, and destructive diagnostics served measurements of several parameters of interest of confined electron clouds. In particular, the single-particle motion of the electrons was investigated and the period of the transverse component of the motion (magnetron rotation) was measured for several configurations of the external fields of the trap. The trends obtained agree well with the results of PIC simulations. In addition, some of the limitations of the present apparatus were determined from measurements of the most important parameters of interest of the trapped electron clouds: maximum achievable particle number and density, transverse size and profile, lifetime, and radial expansion rate.

Given that the conditions envisioned for the plasma lens for LhARA have not yet been simultaneously produced in experiments, it was considered whether a solenoid based on normal-conducting (NC) technology could be used as an alternative to the Gabor lens. A preliminary design was created starting from analytical models, practical recommendations, and conventional values used in the design of accelerator magnets. It was

shown that the design can be adapted to a number of magnitudes of the solenoid current based on the choice of the most suitable current source. Moreover, the size and power consumption of the magnet could be decreased while also achieving better cooling effectiveness by varying the cross-section area of the cable across the layers of the coil. The final designs were verified using 2-D magnetostatic simulations to ensure that they can produce the required field strength and quality. The resulting preliminary designs established approximate values for a few principle factors that need to be taken into account when evaluating the choice of plasma lenses versus NC solenoids: total mass and size of the magnet, required cable length, complexity of the cooling system, and power consumption.

Finally, a study was conducted comparing the beam capture and transport capability of solenoid magnets and plasma lenses in LhARA. The comparison included an evaluation of the number of protons that can be effectively captured from the source in the case of magnetic or plasma-based focusing in a more general setup mostly identical to the front-end of the LhARA beam-lines. Furthermore, the scheme designed to reduce the beam energy spread in LhARA was validated with both solenoids and plasma lenses. These studies involved tracking a realistic distribution of laser-driven protons from the source to one of the end stations of LhARA. To this end, two models for an ideal and a more realistic plasma lens were created and implemented in a couple of particle-tracking codes.

This manuscript is organised in the following way: The first chapter introduce the basic principles of laser-plasma proton and ion acceleration with emphasis on the potential applications of the generated beams and on the main techniques to capture and transport the particles (chapter 1). Then, some general notions regarding charged particle therapy are described and the LhARA facility concept is presented (chapter 2). The basic concepts of non-neutral plasma confinement and proton- and ion-beam focusing are introduced, and the most important previous attempts to build a working plasma lens are discussed (chapter 3).

The central work of my PhD is summarised in the following four chapters, with one chapter each presenting the studies of the anomalous focusing effect produced by a prototype plasma (Gabor) lens (chapter 4), a series of measurements of the parameters of low-density electron clouds confined in a plasma trap (chapter 5), the development and study of a design for a NC solenoid magnet that could act as an alternative to the plasma lens in LhARA (chapter 6), the comparison of the performance of solenoids and plasma lenses in capturing and transporting protons to the in-vitro end station of LhARA (chapter 7).

Finally, the main conclusions of the thesis are summarised and the perspectives for the future development of a plasma lens for LhARA are discussed (chapter 8).

Chapter 1

Novel proton and ion sources

This chapter introduces the basic principles and general aspects of proton and ion acceleration with the use of high-power lasers. First, the fundamentals of modern laser-plasma acceleration techniques are presented (sections 1.1.1 and 1.1.2) and a number of potential applications of the resulting particle sources is discussed (section 1.2). Then, an overview of conventional and novel systems suitable for focusing charged-particle beams is given (sections 1.3.1 to 1.3.4) with emphasis on the capability of these systems to be run at high repetition rates, and to provide strong focusing and large beam acceptance. The chapter concludes with a discussion of the main limitations of the conventional systems investigated so far for capturing and focusing laser-driven protons and ions (section 1.3.5).

1.1 Laser plasma ion acceleration

1.1.1 Acceleration mechanisms

The first experiments reporting proton beams with energies of up to several tens of MeV generated behind thin foils irradiated by intense lasers were reported in 2000 [55–57]. Since then, an extraordinary amount of experimental and theoretical work has been devoted to the demonstration and understanding of the generation of multi-MeV proton and ion beams over a wide range of laser and target parameters. Most of this work has been based on sheath acceleration, often referred to as target normal sheath acceleration, or TNSA, as firstly proposed by [58]. More recent theoretical and experimental studies laid the groundwork for novel acceleration mechanisms such as radiation pressure acceleration [59] or laser breakout afterburner type acceleration [60]. These mechanisms occur in the transparency regime and rely on high-contrast laser systems that are able to irradiate targets as thin as only a few nanometres. However, these alternative mechanisms tend to require much higher laser intensities than are needed for TNSA.

A comprehensive summary of ion acceleration mechanisms from laser-plasma interactions can be found in [22] and [62] and an overview of a few selected regimes that were extensively considered for applications is given in table **??**. As TNSA is the most widely investigated technique that is now known to be robust and effective, it is currently the preferred mechanism for developing applications using relatively low-energy-laser

Table 1.1: A selection of laser-driven ion-acceleration mechanisms and their characteristics in view of potential applications [22, 61]

Laser targets	Laser-plasma regime	Mechanism	Characteristics
Thick ($>1\,\mu$m) foils, solid density	opaque	TNSA	widely studied, broadband energy spectrum, large opening angle at the source, limited maximum ion energies[†],
Thin ($\sim$1–100 nm) foils, solid density, require ultra-high laser contrast	opaque	RPA Light Sail	quasi-monoenergetic spectrum, preferential laser-energy coupling to heavier ion species, possibility of pure ion acceleration
	relativistically transparent	Break-Out Afterburner	quasi-monoenergetic spectrum, preferential laser-energy coupling to heavier ion species
Near critical-density, long plasmas	opaque	RPA Hole-Boring	higher laser-ion efficiency, large energy spread, preferential laser-energy coupling to heavier ion species

[†]with the currently available high-power lasers relative to the other selected mechanisms

driven ion sources at high repetition rates. The main limitation of TNSA stems from the current limits on ion energies that can be produced. The maximum energies observed so far were obtained with Petawatt laser systems at low repetition rate. Increasing the beam energy further to suit possible applications, such as proton radiotherapy, will still require significantly higher laser intensities.

Target Normal Sheath Acceleration (TNSA)

Target normal sheath acceleration is based on the irradiation of thin foils from a few µm to tens of µm thickness by an intense laser pulse in an intensity regime defined by the guideline $I\lambda^2 > 10^{18}\,\mathrm{W/cm^2\mu m^2}$, where I and λ are the laser intensity and wavelength. The basic mechanism is illustrated in fig. 1.1. The initial part of the laser pulse interacts with the front surface of the target, forms a pre-plasma, and transfers some of the laser energy to electrons. The number and temperature of the hot electrons influence the overall acceleration achievable by TNSA. The main pulse propagates through the pre-plasma and accelerates the electrons to MeV-scale energies so that they can propagate into the bulk of the target plasma to the rear surface. A limited number of sufficiently energetic electrons will leave the target, while most of the hot electrons will be held back by the positive space-charge forming within the target. The hot electrons circulate between the front and rear of the target and produce an electron sheath on both sides extending beyond the initially unperturbed surfaces of the target. As a result, this sheath sets up a large accelerating electric field (on the order of $\mathrm{TV/m}$ [63]) which can ionise atoms and rapidly propels the ions out of the foil.

The strength of the sheath fields typically increases for higher electron temperature leading to an enhance-
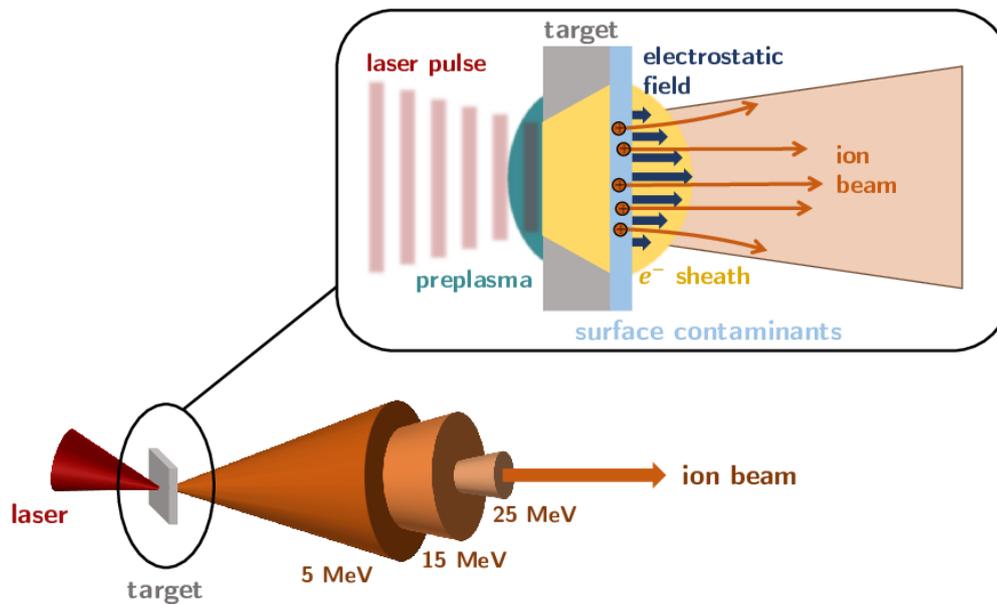
Figure 1.1: Schematic diagram of target normal sheath acceleration (TNSA) from an intense lases impinging on a solid target. Hot electrons from the preplasma form a strong electrostatic sheath field at the rear of the target which accelerates the ions from the target surface into the vacuum. Typical surface contaminants are hydrocarbons and water vapour. The diagram also shows schematically that TNSA generates the more energetic ions within smaller solid angles compared to the slower ions.

ment of the maximum ion energy. A similar improvement can be obtained from the hot electrons recirculating through the target during the acceleration process, under the right combination of target thickness and pulse duration [64]. TNSA from the front target surface is less efficient due to the presence of the preplasma. In principle, TNSA can accelerate any ion species that is present in the surface layers of the target. In most of the experiments, TNSA was observed to preferentially accelerate light ions (protons, carbon and oxygen ions) from contaminant layers rather than ions from the bulk of the target. Protons are always accelerated first as they have the highest charge-to-mass ratio. Hence, contaminants on the target surface must be removed prior to the laser irradiation in order to accelerate preferentially ion species heavier than protons.

1.1.2 Properties of laser-plasma ion sources

Part of the motivation for the extensive research on laser-accelerated ion beams stems from their exceptional properties. First of all, a plasma can support very large electric fields, up to TV/m, via the local charge separation in the target induced by the laser pulse. By comparison, conventional ion accelerators like synchrotrons or linacs are based on accelerating cavities with gradients of up $\sim 50\,MV/m$. As a result, the laser-driven acceleration takes place over very short distances. Even when considering the overall footprint of the laser and target systems, the laser-driven source is considerably more compact than conventional accelerators that can achieve ion energies of several tens of MeV or above.

A second benefit of laser-driven protons is the extremely high number of particles produced per laser shot which is typically between 10^{11} and 10^{13}. Towards the high-energy end of the proton distribution, this number can drop as low as 10^7–$10^8\,particles/MeV/sr$ [63]. Furthermore, the ion bunches have an ultra-short length
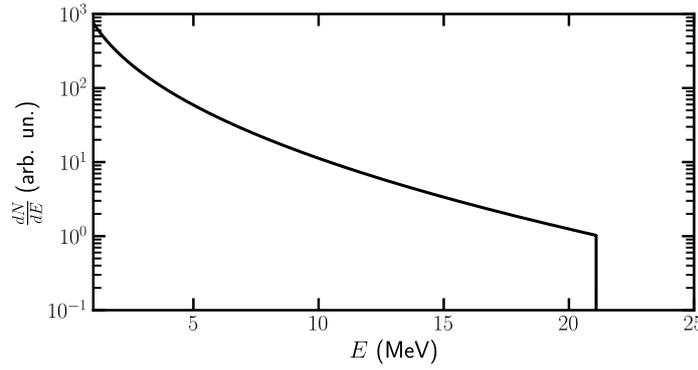
Figure 1.2: The main features of a typical energy spectrum of protons produced by TNSA. The particle number decreases with an increase in the kinetic energy E according to a pseudo-thermal spectrum up to a maximum cut-off energy. The spectrum is calculated from the well-known collisionless plasma expansion model proposed in [67, 68] by tuning one of the plasma parameters, the electron temperature T_e, to give a cut-off energy around $20\,\mathrm{MeV}$.

at the source dictated by the laser pulse duration, typically of the order of ps or shorter, and the source is by construction pulsed and triggerable. The protons and ions exit the target being accompanied by a cloud of electrons, which leads to space-charge neutralisation until the two species separate. Finally, the beams are extremely laminar with very low emittance when compared to conventional accelerators, less than $0.004\,\pi\,\mathrm{mm}\cdot\mathrm{mrad}$ and $10^{-4}\,\mathrm{eV}\cdot\mathrm{s}$ in the transverse and longitudinal phase-spaces, respectively [65].

However, with the currently available high intensity lasers, the maximum proton and carbon ion energies that have been achieved via TNSA are still much below the energies that are routinely reached at conventional accelerator facilities for high-energy physics or carbon ion radiation therapy, for example. A further difference is the energy spread of the beams produced. While conventional accelerators provide beams with extremely low energy spread, laser-driven sheath acceleration results in beam with a broad and continuous energy distribution. Significant research effort is currently invested in developing new acceleration techniques that can provide spectral shaping of the beam from the source. Alternatively, the conditioning of the beam energy spread for possible applications can be done with the use of the beam transport line based, for example, on longitudinal phase rotation by RF cavities and on energy collimation due to energy-dependent beam-focusing [66]. In view of the possible ground-breaking applications of laser-driven ion beams from TNSA, some of the beam properties at the source act as drawbacks. These properties are described in more detail below.

Energy spectrum

One of the properties of proton and ion beams produced by the interaction between a high-intensity laser and a solid target is the typical broad energy spectrum which is continuous up to tens of MeV (see, for example, [57, 69, 70] and [71]). The broad energy distribution is caused by the acceleration of particles from different depth within the target. Both protons and heavier ions are accelerated from water vapour or hydrocarbon contamination that are always present on the back surface of the target due to the limited achievable vacuum conditions. As the thickness of the contamination layer increases, the resulting energy spectrum becomes broader.

There are other various laser and target parameters which dictate the resulting particle energy spectrum. For proton acceleration via the TNSA scheme, some of the parameters that have been experimentally confirmed to influence the energy spectrum are: the material of the target [72], the target thickness [64], laser intensity and pulse energy [22], and the presence of a preplasma [73].

A number of models have been developed to try and explain the TNSA mechanism and predict the main features of the characteristic energy spectrum. One of the most well known models, given in [67, 68], uses a collisionless plasma expansion model to predict the energy spectrum based on the assumption that the electron distribution corresponds to a Boltzmann distribution. Firstly, the model allows the prediction of the maximum energy at which the number of ions drops off sharply. Secondly, the kinetic energy (E) distribution is obtained as an exponentially decreasing spectrum:

$$\frac{dN}{dE} \propto \frac{1}{\sqrt{Zk_BT_eE}} \exp\left(-\sqrt{2E/Zk_BT_e}\right), \tag{1.1}$$

where Z is the charge number of the accelerated ion species. One limitation of the model is that it requires a priori knowledge of the electron temperature T_e. The main features of a typical proton energy spectrum obtained from TNSA is shown in fig. 1.2.

With recent experimental advancements, proton beams with cut-off energies in excess of $85\,\mathrm{MeV}$ have been obtained by TNSA from submicrometer thick plastic targets [74]. Other less well-investigated acceleration mechanisms can lead to higher proton energies. One examples is a hybrid scheme that relies on RPA at the target front and TNSA at the rear [75].

For heavier ions generated by laser-plasma acceleration, the energy spectrum shows the same exponential decrease as in the case of protons [76]. Medium energy femtosecond class lasers were previously used to investigate the acceleration of carbon ions and generate maximum ion energies of approximately $5\,\mathrm{MeV/u}$ for relatively thin ($\sim 100\,\mathrm{nm}$) solid targets. Higher maximum energies in excess of $30\,\mathrm{MeV/u}$ were obtained from femtosecond class lasers with the use of ultra-thin ($\sim 10\,\mathrm{nm}$) or difficult to manufacture targets to accelerate carbon ions selectively [77, 78]. In the TNSA regime, target surface cleaning techniques were shown to enable the acceleration of heavier ions including carbon and oxygen, to maximum energies of a few $\mathrm{MeV/u}$ [79, 80]. Proof-of-principle experiments aimed at the generation of oxygen ions have been carried out using the same technique of controlling and removing contaminants from targets that can be used at high repetition rates [81]. One of the highest achieved energies for carbon ions ($>80\,\mathrm{MeV/u}$) was reported by [82].

Beam opening angle at the source

In addition to the large energy spread, proton and ion beams produced by a laser-driven source are generally known to have large opening angles. Values for the full beam-envelope opening angle of at least a few tens of degrees have been reported for proton beams [83]. Typically, this angle decreases when the proton energy increases as schematically illustrated in fig. 1.1. In other words, the flux of protons accelerated to higher kinetic

energy is contained by a smaller solid angle compared to the flux of protons that propagate at lower kinetic energy. For example, in fig. 1.1, the protons with kinetic energy of $25\,\mathrm{MeV}$ travel closer to the central axis at a certain distance from the target compared to the protons with kinetic energy of $5\,\mathrm{MeV}$ that reach further away from the axis at the same position along the beam-line.

A comparison of the energy dependence of the opening angle of the beam at the source is presented in [65] for three different laser systems. The particles with the highest energies are emitted at half-angles of approximately $5°$. At about 30% of the maximum energy, the opening angle reaches a maximum and stays constant for lower energies. It is important to make the distinction between the 'opening angle' at the source and the beam divergence, as noted in [65]. The divergence of the protons increase slightly with energy, while the size of the emitting area decreases with proton energy. As a result, the opening angle measured experimentally decreases with energy.

While the broad energy spectrum and the wide opening angle at the source present challenges for any technique that can efficiently capture laser-driven ions, the relatively high energy of the particles at the source is an advantage when compared to conventional ion sources. A brief overview is given below of the influence of the particle energy on the evolution of the transverse beam size.

Space-charge effects

A particle of charge q that travels inside an unbunched beam of particles of the same charge is pushed outwards by the electrostatic Coulomb repulsion. At the same time, each particle travelling with a velocity $v = \beta_0 c$ represents an electric current that creates a magnetic field. This results a magnetic force that is radial, attractive, and acts on each particle in a travelling beam. For a highly parallel beam (paraxial approximation) with uniform current density and cylindrical symmetry, the sum of the electrostatic and magnetic forces has an overall repulsive effect given by [84]

$$F_r = qE_r(1 - \beta_0^2) = \frac{qE_r}{\gamma_0^2}, \tag{1.2}$$

where E_r is the radial electric field produced by the charge density in the beam and $\gamma_0^2 = 1/(1 - \beta_0^2)$. The ratio of the magnetic to electric force is β_0^2 and is independent of the beam distribution. Moreover, at relativistic energies, the beam electric force is almost balanced by the magnetic force. For non-relativistic particles (like low-energy ion beams), the magnetic component is negligible and the overall effect of the beam space-charge is defocusing.

In a drift space, without any external focusing fields, an ion beam will expand due to its space charge. If the beam is paraxial, uniform, and has cylindrical symmetry, the radius of the beam, r_{env}, evolves along the drift space as [84]

$$\frac{d^2 r_{\mathrm{env}}}{dz^2} = \frac{K}{r_{\mathrm{env}}}. \tag{1.3}$$
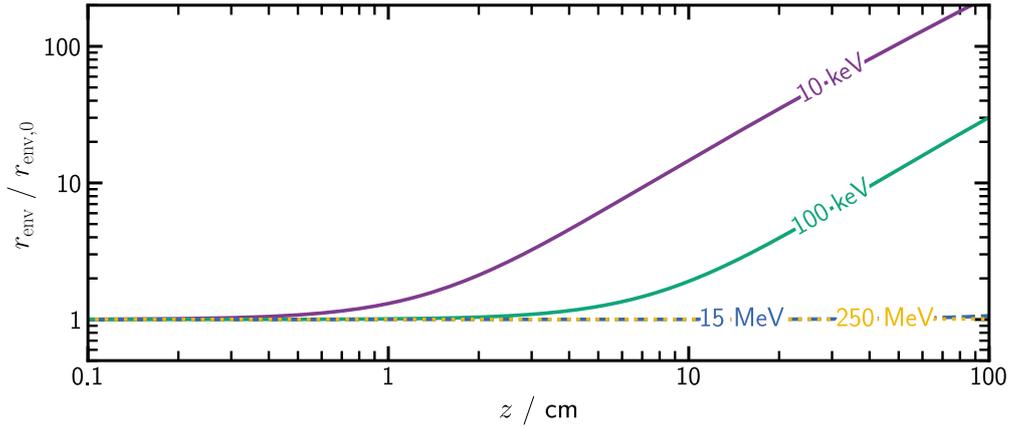
Figure 1.3: The envelope radius, r_{env}, of a proton beam that propagates through a field-free region for a few relevant kinetic energies. The beam has a $10\,mA$ current and a radius at the source $r_{env,0} = 1\,mm$. The high-energy beams propagate across long distances with relatively small increase in the transverse beam size due to space-charge effects.

The generalised perveance, K, is a dimensionless parameter that defines the magnitude of the space-charge effects on the beam:

$$K = \frac{qI}{2\pi\epsilon_0 mc^3 \beta_0^3 \gamma_0^3}. \tag{1.4}$$

One interpretation of K is as the ratio between the strength of the space-charge repulsion in the rest-frame of the beam, which is proportional to the beam current I, and the effect of the relativistic time dilation when transforming to the laboratory frame, which is more significant at higher energies.

It is important to note that realistic ion beams are not uniform in the transverse plane and, hence, the space-charge effects are fundamentally non-linear. Due to the non-linear force, the emittance of the beam increases as it is transported along an accelerator beam line. Since the transverse acceptance of any transport beam line or acceleration stage is limited, there is a constraint on the maximum current from the source that can be successfully injected into the rest of the accelerator. This constraint strongly depends on the strength of the space-charge effects starting from the particle source.

An example of the impact of the space-charge effects on the transverse size of a proton beam is shown in fig. 1.3. At low energy, below $\sim100\,keV$, the beam expands radially very quickly and most of the beam is lost in the absence of any focusing fields, space-charge compensation scheme, or post-acceleration. Conventional ion sources suffer from this space-charge limitation on the maximum current that can be extracted given that the first stage of acceleration immediately after the source brings the beam to an energy of up to $\sim60\,keV$ (see table 1.3). In contrast to conventional sources, laser plasma acceleration typically generates proton beams with an energy from a few MeV to several tens of MeV. Within this energy interval and above, the space-charge effects are much less significant compared to the case of beams from conventional ion sources. As a result, laser-driven ion sources are much more suitable to applications that require high beam currents or the delivery of a large number of particles per bunch.

1.2　Potential applications of laser-driven ion beams

A significant number of applications and potential applications have been found for laser-driven ion sources given the unique properties of the beams that they generate. Some applications can already be implemented with the current beam parameters: proton radiography [85], warm dense matter generation [86], pulsed radiobiology [23], compact sources of neutrons [87] or medical isotopes [88], ion implantation [89]. Other proposed uses of laser-driven ion sources require further advancements in their performance: fast ignition in energy production by inertial confinement fusion (ICF) [90] or ion-beam cancer therapy [91]. A few of these applications with a societal impact are described in more detail below. These are selected based on their potential to benefit from high-intensity ultra-short bunches, high repetition rate beam delivery, and monoenergetic ions, and thus to benefit from an efficient beam-capture and -transport system. A more comprehensive review of the many applications of laser-driven proton and ion sources can be found in [22].

At present, the repetition rate of ultra-high-intensity lasers is usually $10\,\mathrm{Hz}$ or less. With the development of a target system that is capable of sustained operation [92], the overall source system may achieve operation at a comparable repetition rate. More advanced targetry, such as liquid-sheet targets [93] or cryogenic hydrogen jets [94], may allow operation at higher repetition rates. At $10\,\mathrm{Hz}$, 10^{12} protons/pulse are required within the energy range of interest to give an average current of $\sim 1\,\mathrm{\mu A}$ which is comparable to a synchrotron. Although the stability of current laser-driven proton sources is not sufficient for the industrial and medical applications, work is ongoing to develop active optimisation and stabilisation techniques [95].

Ion implantation

Ion implantation has been identified [96] as one potential application of high repetition rate ion sources in the moderate proton energy range of around $1\,\mathrm{MeV}$. Some of the advantages include the delivery of higher charge states that can be implanted to a greater depth and the capability to provide a wide variety of ion species. The most important challenges are that the proton or ion source should satisfy the requirements of a specific ion yield within a fixed time and, correspondingly, to provide a specific average ion current. A simple evaluation of a $0.1\,\mathrm{J}$, $1\,\mathrm{kHz}$ laser-system [22] shows that in order to ensure the required ion fluence, the ion source must provide an average current of $0.02\,\mathrm{mA}$, which is about an order of magnitude lower than that of a conventional machine [96]. One assumption is that all the ions can be focused into an area of $1\,\mathrm{cm}^2$. However, lower ion yields are required if the beam can be focused to a smaller area. If the presently available ion yield can be increased to make laser-driven ion sources a potential candidate for industrial applications, the efficient capture and transport of the low energy ions from the source is critical.

Radioisotope production

At slightly higher beam energies up to several MeV, protons can drive nuclear reactions [97] and produce useful isotopes [88, 98]. The figure of merit is the yield per unit time for protons with energy larger than the threshold

of a specific nuclear reaction. Radioisotopes play an essential role in diagnosis and therapy in nuclear medicine, with an annual increase in the number of procedures performed every year based on them [99].

In one of the major diagnostic techniques, positron emission tomography (PET), the patient is injected with a pharmaceutical labelled with a short-lived positron-emitting isotope. The radioisotopes used in PET are typically ^{11}C, ^{13}N, ^{15}O and ^{18}F. The nuclear reactions employed to produce these isotopes are ignited by $20\,\mathrm{MeV}$ protons or similar energy deuterons from cyclotrons. The short half-life of the PET radioisotopes limits their use to certain physical distances around the cyclotron facility. Moreover, accelerating the driver beams comes with the associated issues of large size and cost, and the need for extensive radiation shielding, all which might be addressed by high-repetition-rate laser sources. Without any optimisation of the laser-target system, a table-top laser generating $1\,\mathrm{J}$ pulses at $10\,\mathrm{Hz}$ could reach an activity of $150\,\mathrm{MBq}$ [100] which is only a few times lower than the activity necessary for patients.

The most widely used radioisotope is ^{99m}Tc which accounts for more than 80% of all nuclear medicine procedures and diagnostic scans [99]. ^{99m}Tc has a longer half-life of $6\,\mathrm{h}$ and is typically produced by the spontaneous decay of ^{99}Mo which is generated by the fission of uranium in nuclear reactors. An alternative is to produce ^{99m}Tc by using a cyclotron. However, a laser-driven proton source may provide a significantly more cost-effective solution. Both the production of ^{99}Mo and ^{99m}Tc were demonstrated and studied at several laser-plasma accelerator facilities. A summary of these efforts is provided in [99]. The overall picture of the production capability of different laser facilities indicates that an increase in the laser repetition rate is required for the radioisotope yields to meet the dose requirements for clinical use. In parallel, advanced accelerator beam-line lattices have been proposed to increase the yield from laser-driven protons by actively manipulating the particle distribution from the source [66].

Radiobiology

There are two main motivations that led several research groups to initiate experimental activities in which cellular samples were irradiated by laser-driven proton beams. Firstly, one aim is to identify and better understand the challenges of using the laser-plasma interaction as a proposed compact ion source for future cancer therapy facilities. Secondly, the ultra-short duration of the proton bursts and the high particle number per bunch enable the irradiation of cells and tissues at ultra-high dose rates, and, thus, to access unexplored regimes of radiobiology. Furthermore, with appropriate modifications of the target system, different ion species, such as carbon, can be used for irradiation from the same source, removing the need for multiple ion sources being installed in the same facility.

In a typical arrangement, doses of up to a few Gy can be delivered to the cells in short ion bunches of $\sim$ns duration. While conventional RF accelerators deliver typical average dose rates of $\sim 0.1\,\mathrm{Gy/s}$, cells were irradiated with laser-driven protons resulting in single-shot on-cell dose rates on the order of $10^7\,\mathrm{Gy/s}$ [25, 101] up to $10^9\,\mathrm{Gy/s}$ [102, 103]. Given this enormous difference in dose rate, experiments initially have focused on assessing the biological effect of laser-driven ions in comparison to conventional ion or photon beams. To this end,

several experiments [101, 102, 104] measured the relative biological effectiveness (RBE) of laser-accelerated protons to be comparable to that of conventional radiation. The comparison was based on standard cellular assays and the evaluation of either the damage inflicted to the DNA of the cells or the fraction of cells surviving the irradiation. Further therapeutic potential of laser-driven protons has been identified by experiments measuring the long-term effects on cells irradiated at low levels of dose, the so-called sub-lethal effects. Reports [105, 106] indicate a less negative impact of laser-driven protons on the cells compared to irradiation with conventional proton beams mainly due to the ultra-short proton pulses originating from laser-plasma acceleration and, in turn, the temporal distribution of DNA-damaging events that may modify cellular and tissue radio-sensitivity. In a clinical context, this could lead to lower damage to the healthy cells traversed by the ions on their way to a tumour.

Due to this huge potential of laser-driven ions beams to provide a step forward in the provision of local tumour control whilst sparing normal tissue, several laboratories are engaged in radiobiological research with ultra-high dose rate protons and pilot studies have recently started on small-animal models [27]. Efficient capture, transport and energy-filtering of the laser-generated protons is crucial in order to allow the delivery of significant doses at energies above the current limit of $\sim 30\,\mathrm{MeV}$. Radiobiological investigations [23] have been so far below this limit where comparison to RF-accelerated proton beams cannot be accurately achieved due to a lack of robust dosimetry at such low energies. Furthermore, beam-transport and focusing techniques are required to provide uniform irradiation of a large surface area of cells.

Cancer therapy

In the context of hadron-beam cancer therapy, laser-driven sources were proposed soon after the discovery of TNSA as highly efficient ion injectors for a conventional post-acceleration stage [107] or as a complete replacement of a conventional medical accelerator [19–21]. The main motivation behind these proposals is that laser-driven sources may lead to advantages in terms of the footprint of the ion accelerator or the overall treatment facility and to lower cost. Moreover, it is also hoped that laser-driven sources will provide the path to more tunable medical accelerators which could rapidly change the characteristics of each particle bunch by varying the parameters of the laser pulse from shot to shot.

In order to penetrate into the body sufficiently, the energy window used in proton beam therapy ranges between 60 and $250\,\mathrm{MeV}$, depending on the location of the tumour. Since the location where most of the radiation dose is deposited in the tissue depends on the energy of the beam, optimal dose delivery over the tumour region requires small deviations of the particle energy from the nominal value: $\Delta E/E \approx 1\%$. With the recent experimental results reporting proton energies of around $100\,\mathrm{MeV}$, the lower energy threshold for medical applications has been achieved. However, the beam energy spread from the source still needs to be significantly reduced. Thus, the complexity and size of the magnetic devices that were proposed up to now for shaping and transporting the beam can exceed those of a conventional accelerator and reduce the advantages of laser-driven sources. For example, one method for the monochromatisation of a laser-driven proton beam is

to produce a rotation in the longitudinal phase-space with a carefully synchronised RF cavity [108]. While the phase-rotation technique can be applied at high-repetition rates, several meters of RF cavities are required to potentially achieve an energy bandwidths of a few MeV for a $\sim 200\,MeV$ proton beam. Alternatively, the need for external devices may be eliminated by the development of new regimes of proton and ion acceleration as well as special target designs that can control the beam divergence and spectrum directly at the source.

Proton and ion accelerators for hadron therapy are required to provide good control of the energy of the particles and the irradiation dose that are delivered to the patient. Future therapy facilities based on laser-driven ion source may achieve the required control over the beam parameters using online beam monitoring and beam-lines that are capable to select and deliver particles within a narrow energy interval.

Fast ignition inertial confinement fusion

Inertial confinement fusion (ICF) is a proposed fusion scheme that employs energy deposition on a fuel capsule to compress and heat the fuel and, thus, initiate the nuclear fusion reactions. In the classical scheme, a laser is used to irradiate the inside of a metal cylinder, with laser energy being converted into X-rays which hit the outer layer of the fuel capsule located at the centre of the cylinder. As a result, the outer layer expands outward and produces a compression force and powerful shock waves in the core of the target. The concept of Fast Ignition [109] was proposed as a means to increase the gain, reduce the driver energy and relax the symmetry requirements for compression by pre-compressing the cold fuel and subsequently ignite it with a separate high-intensity laser or a short pulse of particles. The use of laser-driven ions as a trigger was proposed [90] to take advantage of their capability to maintain almost a straight trajectory while traversing the compressed target, their excellent coupling efficiency to the fuel, their energy-deposition profile highly concentrated towards the stopping point (Bragg peak), and their focusability.

To achieve stopping at the right range in the fuel, the required particle energies for fast ignition are 5–$15\,MeV$ for protons and 400–$500\,MeV$ for carbon ions [110]. Quasi-monoenergetic beams are preferred due to their better coupling with the compressed fuel (better localised Bragg peak). Given this relatively low energy per particle, a very large number of protons ($>10^{16}$) or ions ($>10^{14}$) is required to deliver the necessary energy of $\sim 10\,kJ$ for ignition [111]. As in the case of isotope production, the laser-ion conversion efficiency is among the most important parameters for ion-driven fast ignition, too. As an approximate threshold, a 10% laser-to-ion conversion efficiency would start to bring a significant reduction in the required laser energy compared with the traditional laser-induced fusion. Furthermore, probably the most challenging condition of fast ignition is the requirement for the ion beam diameter not to exceed $\sim 40\,\mu m$ at the location of the dense core [112].

The non-monoenergetic nature of the laser-driven ion spectrum causes significant time-of-flight spread in the beam over long distances. To deliver a sufficiently short (~ 10s of ps) ion bunch to the target, the ion source must be situated close ($\sim 1\,mm$) to the target which, in turn, requires re-entrant cones and protection foils as part of a complicated target design. In principle, one alternative is to capture the ion flux efficiently from the laser-driven source, filter the energy spectrum, apply RF bunch compression and focus the beam on the ICF

Table 1.2: The beam requirements for a selection of applications of laser-driven ion sources. Values taken from [22], [62] and refrences therein.

Application	Key beam requirements
Ion implantation	high rep. rate and high flux ($\sim100\,\mu\text{A}$), $<1\,\text{cm}^2$ beam-size, $E_k \approx 1\,\text{MeV}$
Radioisotopes production for PET	high rep. rate or high flux (a few hundred μA), $E_k \approx 6\,\text{MeV}$
Fast ignition inertial confinement fusion	$\Delta E/E \approx 10\%$, $E_k \approx 5\text{–}15\,\text{MeV}$ (protons), $E_k \approx 400\text{–}500\,\text{MeV}$ (carbon), $\sim10\,\text{kJ}$ proton or ion beam ($>10\%$ laser-ion efficiency), beam-energy deposition in 'hot spot' volume $\sim25\,\mu\text{m}^3$ within $\sim20\text{–}50\,\text{ps}$
Ion radiotherapy	high rep. rate, $\Delta E/E < 1\%$, $E_k \approx 250\,\text{MeV}$ (protons), $E_k \approx 400\,\text{MeV/u}$ (carbon), $\sim(1\text{–}5) \times 10^{10}$ particles/s

target with the use of a beam-transport line. Moreover, it has been shown in simulations that good control over the location of the focusing point of the ion beam relative to the compressed core is crucial to reduce the ignition energies [112]. Further beam-focusing capabilities that can produce ring-shaped proton beams may result in a reduction of the energy requirements in a scenario based on a combination of two spatially shaped proton beams [113].

Fundamental high-energy physics research

Apart from the industrial or medical applications, laser-driven sources were also proposed as injectors of high power ion beams for large-scale basic research facilities. The advantage over conventional accelerator-based injectors is given by the lower cost and the more compact size. One direct application for sufficiently high-energy (hundreds of MeV or GeV) laser-driven proton beams is the production of secondary particles. For example, the threshold for pion production corresponds to a $\sim140\,\text{MeV}$ proton beam [114], but the production cross-section increases at higher proton energies. At the higher threshold of $\sim1\text{–}2\,\text{GeV}$, proton beams can generate pions of which decay results in a short burst of neutrinos [115]. This non-conventional source for low-energy neutrinos was proposed in close synergy with the use of GeV proton driver beams for neutron spallation sources. To make full use of the foreseen increase in beam energy from laser-driven proton sources, a key requirement is the capability to capture the particles from the source efficiently and to transport them with low losses towards the next stages of acceleration, or, potentially, directly focus them into the interaction region. Ideally, the reduced cost and size of the source must be matched by a similarly compact and cost-effective capture and focusing system that is suitable for high proton-beam energies.

Most of the applications named above can only be realised when the uniquely high intensity of the laser-driven bunches is preserved while simultaneously gaining control over their angular divergence and energy distribution. Table 1.2 summarises the principal beam requirements for a number of applications of laser-driven

Table 1.3: Main features of Low Energy Beam Transport (LEBT) lines for high-intensity H$^-$ sources.

Ion source	LEBT layout	Length† (m)	Extraction voltage (kV)
SNS [118]	two electrostatic lenses	0.12	65
CERN [119]	electrostatic lens and two solenoids	2.10	45
JPARC [120]	two solenoids	0.66	50
CSNS [121]	three solenoids	1.65	50

†from the plasma extraction plane to the entrance of the first acceleration stage

ion sources. The control system that can provide some of these requirements must be compatible with source operation at a high-repetition rate and its design must be driven by taking into account the cost, size and reliability of the facility. The following section describes several conventional and novel techniques that provide efficient focusing and transport of ion beams. The overall efficiency in providing the required control over the beam transport is discussed for each focusing system in the context of bridging the gap between the unique properties of the laser-driven ion sources and their potential applications.

1.3 Techniques for ion beam capture and transport

For the production of beams with high brightness it is important to preserve a small emittance in the low energy beam transport (LEBT) system that is used to transport and match the ion beams from the source to the first stage of acceleration. The ion beams produced by conventional plasma sources, such as multicusp volume or surface plasma sources, are typically divergent and must be refocused before entering the first acceleration stage, usually an RFQ. To maximise the transmission through the accelerating structure a certain combination of beam size and convergence is needed at the entrance plane.

Facilities that produce high-intensity proton beams rely on electrostatic or magnetic focusing. Reliable operation of an electrostatic LEBT has been demonstrated with H$^-$ beams up to $60\,\mathrm{mA}$ [116]. Higher current ($>100\,\mathrm{mA}$) proton beams have been successfully transported with a magnetic LEBT [116, 117]. A summary of the LEBT specifications for a selection of accelerator projects is given in Table 1.3. Due to the low energy of the beams produced by conventional ion sources (up to $\sim 60\,\mathrm{keV}$) space-charge forces are particularly strong for high intensity beams. The electric repulsion between the ions tends to defocus the beam and leads to emittance growth due to its strongly nonlinear nature. The defocusing effect of the space charge can be compensated by transporting the low energy beam through a heavy gas in the so-called space charge compensation (SCC) regime. However, partial SCC and the collisions with the gas molecules can cause large beam losses and degradation of the beam brightness.

For ion beams produced by laser-driven sources, the space-charge defocusing effect becomes significantly less impactful on the beam transport due to the higher energies at which the particles are produced, typically a

few MeV and up to several tens of MeV. However, the beam divergence at the source is larger for the laser-driven mechanisms compared to conventional plasma sources. To transport the highly divergent beam efficiently, the focusing elements downstream of the source must have large focusing strengths and large apertures. Furthermore, a compact beam-line requires focusing elements with short focal length which at higher beam energy can be achieved with large electric or magnetic fields, thus increasing the complexity of the beam-line.

Laser accelerated protons have been successfully captured and collimated using a variety of techniques. For example, laser-driven micro-lenses [122] were demonstrated to focus protons and select an energy interval of the beam by irradiating a micro-cylinder with a short high-intensity laser pulse. A drawback of this setup is its complexity since two synchronised high intensity laser beams are required. Another recently developed scheme relies on a miniature helical coil positioned at the rear of the target [123] that also post-accelerates the guided protons. Highly directional proton beams were obtained with a narrow energy spread [124] due to the strong electromagnetic pulse that travels along the coil at the same speed as the protons. However, this technique is destructive and cannot yet be implemented at high repetition rates, as well as suffering from a lack of pointing stability.

In addition to these innovative techniques, proton beams were successfully captured and focused by devices that may serve in future beam lines for radiobiological studies or even patient treatment. Successful systems for capture, transport, and focusing of protons so far have been based on: permanent magnet quadrupoles [30–32]; a quadrupole-triplet lens [28, 29]; high-field pulsed solenoids [26, 33–35]; and active plasma lenses (APL) [25].

In this section, several focusing devices are presented including a discussion of their suitability to work in combination with laser-driven sources. The section starts with the focusing elements that are conventionally used in LEBT lines and then continues with systems that have been designed more recently to capture and focus laser-driven particle beams.

1.3.1   Electrostatic focusing

Electrostatic fields are used for focusing immediately downstream of particle sources (injectors) and for transport of ions at low energies. In an electrostatic LEBT, the compensating particles are quickly removed by the electric fields and, thus, the beam is transported without any SCC. The beam lines can be very compact since there is no need for a repelling electrode to trap the compensating particles. However, beam losses can lead to high voltage breakdowns and beam trips.

For given maximum strengths of the electric ($\mathbf{E}$) and magnetic ($\mathbf{B}$) fields that can be achieved in practice, the electrostatic component of the Lorentz force, $\mathbf{F} = q\left(\mathbf{E} + c\boldsymbol{\beta} \times \mathbf{B}\right)$, will be stronger than the magnetic component at sufficiently low β. While a precise determination of the threshold β at which the forces are equal depends on the type and size of the focusing lens, an order of magnitude estimation of it can be made by using typical attainable field strengths [125]. For reliable operation, the electric fields in vacuum are limited to about 10^7 V/m for distances less than 10^{-2} m; beyond this field, electrical breakdown in vacuum occurs [125]. For an iron dominated quadrupole magnet with a 1 T field at the pole tip and an electric field of 10^7 V/m, the forces are
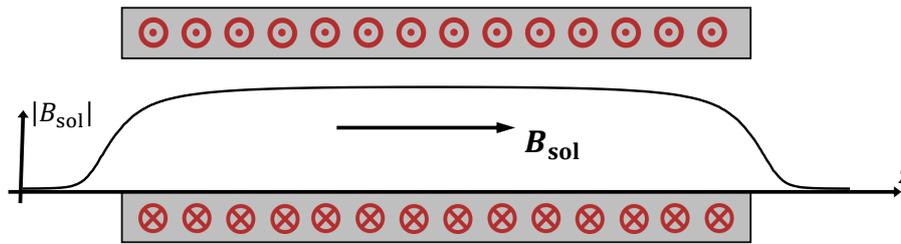
Figure 1.4: Schematic diagram of a solenoid magnet. The longitudinal magnetic field is produced parallel to the *z*-axis by a cylindrical current-carrying coil. The strength of the field, $|B_{sol}|$, decreases across a finite distance at the two ends of the coil defining the fringe-field regions.

equal at $\beta = 1/30$ ($\sim$520 keV protons). Thus, at beam energies above 1 MeV which are typical to laser-driven sources, magnetic focusing becomes more efficient over electrostatic focusing. In conclusion, electrostatic lenses are very efficient in terms of fabrication, compactness, and operational cost, but they are strongly limited to low energy ions.

1.3.2 Magnetic focusing

There are two types of magnetic lens conventionally employed in accelerator beam-lines. Solenoids (fig. 1.4) are the most common focusing and transport systems, especially for low energy particle beams [126]. For focusing high energy beams, magnetic quadrupoles (fig. 1.5) are more common. These two types of resistive focusing magnets are capable of generating magnetic fields of a few Tesla and operate in either AC or DC mode. Under the thin lens approximation, the focal lengths of a quadrupole and a solenoid scale inversely proportional to the momentum and momentum squared, respectively, of the ion species in the beam.

Magnetic quadrupole lenses

The quadrupole gradient $\partial B_y/\partial x$ is typically used to characterise the strength of a quadrupole magnet and it reduces to the maximum achievable field strength divided by the radius of the bore, B_{quad}/r. Because the quadrupole focuses in one transverse plane and defocuses in the other, two quadrupoles separated by a length L and rotated by $90°$ with respect to one another are needed to form a doublet that achieves focusing in both planes. A quadruple doublet has a total focal length of $1/f_{doublet} = L/f_{quad}^2$ [127]. Therefore, both the solenoid and the double-quadrupole systems have a focal length proportional to $p^2/(ZB)^2$. Considering a quadrupole doublet as a capture system for ions produced by the TNSA, the defocusing of the first quadrupole increases significantly the already large divergence of the beam in one of the transverse directions.

Permanent Halbach magnet assemblies have been successfully applied and tested as compact quadrupole focusing systems. However, the use of permanent magnet quadrupoles (PMQs) was observed to lead to high particle losses. Large magnetic field gradients of several hundreds of T/m are required to capture protons at an energy above 10 MeV. Therefore, high field gradient is achieved by quadrupole designs with small apertures typically of only a few millimetres, which results in significantly low transmission. For example, a PMQ doublet
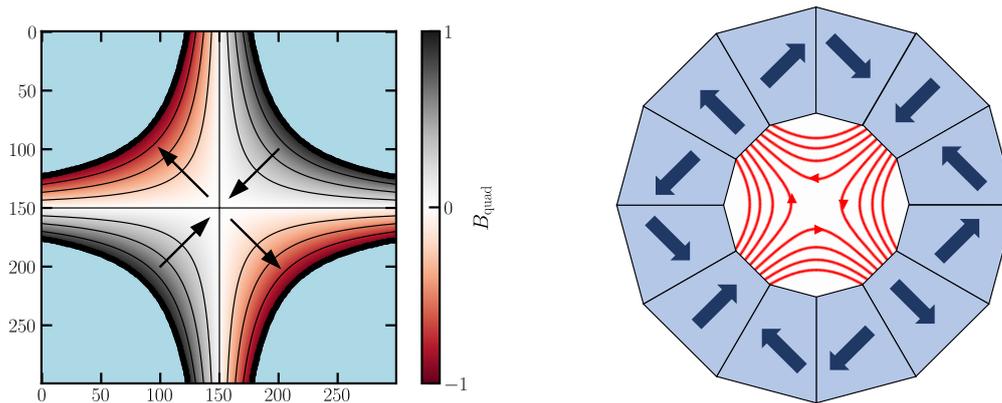
Figure 1.5: **Left:** The transverse magnetic field map of an electromagnetic quadrupole. The field strength is normalised to its maximum values attained near the four poles of the magnet. The black arrows indicate the direction of the force on a charged particle travelling orthogonal to the page. **Right:** Schematic view of a 12 segment Halbach design of a permanent magnet quadrupole (PMQ). The thick arrows indicate the direction of the magnetisation while the curved arrows show the direction of the magnetic field lines.

with field gradients of $55\,\text{T/m}$ and $60\,\text{T/m}$ was demonstrated to focus $(2.4 \pm 0.1)\,\text{MeV}$ protons with less than 3% capture efficiency [31]. Higher energy, $(14 \pm 1)\,\text{MeV}$, protons were focused with the use of smaller bore PMQs with higher gradient around $500\,\text{T/m}$ [30]. The reduction in the size of the bore, however, also led to a capture efficiency of only 0.1%.

In a PMQ doublet, the particle loss is also attributed to the small effective acceptance of the second PMQ as the first PMQ defocuses the beam in the orthogonal transverse direction [31]. Furthermore, the spectrum transmitted by PMQs is not perfectly selected and it requires an additional energy selection system to reduce the energy width of the beam [32]. In the view of future radiotherapy facilities that may generate $250\,\text{MeV}$ proton beams for treatment, the limited magnetic field of PMQs (on the order of $1\,\text{T}$) make them unsuitable to capture the highly divergent protons with sufficient efficiency. The current permanent magnet technology also offers limited tunability of the focusing strength to allow fast, pulse-to-pulse, changes in the delivered beam energy.

An electromagnetic quadrupole triplet lens has been successfully used to capture and focus laser-driven protons with a simulated collection efficiency of the lens of about 88% for the particles within $\pm 50\,\text{mrad}$ [28]. For increasing beam energies, above a few MeV, the advantage of the quadrupole doublet/triplet is that the pole-tip fields of these magnets can remain within the limit of iron saturation at room-temperature and still provide a focal length comparable to pulsed or superconducting solenoids [128]. An analytical scaling of the focusing power at higher energies shows that the transmission of a quadrupole triplet is generally lower compared to that of a solenoid [128].

Pulsed high-field solenoid lenses

On top of the larger transport efficiencies due to the symmetric focusing in the two transverse directions, a solenoid lens can accommodate larger acceptance angles compared to the configurations based on magnetic quadrupoles. Resistive solenoid magnets are limited by their DC power supply and Ohmic heating, and, thus, are not suitable for focusing beams with energies above a few MeV. However, non-destructive pulsed solenoids are able to achieve temporarily significantly more intense peak magnetic fields up to several tens of Tesla. These magnets have a significantly smaller size and weight and the magnetic field can be tuned independently per pulse as dictated by the pulsed power supply. One can take advantage of the pulsed power technology because the laser accelerated ions have a bunch length on the picosecond or nanosecond scale.

Good beam transmission of around 50% was measured on a beam-line based on two pulsed high-field solenoid lenses [26]. The two identical solenoids had a $40\,\mathrm{mm}$-diameter bore and were driven by capacitor-based pulse generators to achieve an on-axis magnetic field strength up to approximately $19.5\,\mathrm{T}$. The ms-long current pulse driving the solenoid is long compared to the duration of the laser pulse [36]. However, the maximum repetition rate of the pulse generators is limited to two to three pulses per minute. The limitation was observed during continuous operation at this maximum frequency which leads to Ohmic heating and an increase in the temperature of the solenoid. A higher operation frequency and temperature of the coil could lead to outgassing from the compounds and adhesive used to reinforce the coil and could enable sparking between the windings.

Further limitations were reported on the maximum current that can drive a pulsed solenoid. As the capture solenoid has to be placed close to the target to increase the angular acceptance, the plasma that is produced when the laser hits the target expands and penetrates into the solenoid creating a short circuit inside the coil [33]. In addition, the close proximity of the solenoid to a metallic target needed further investigation. The non-zero magnetic field at the location of the target and the strong Eddy currents induced in the metal can lead to a bending of the target as reported in [33] by several degrees. As the target is moved out of the focal plane of the laser by more than the target thickness (20–$50\,\mathrm{\mu m}$) the laser intensity at the target drops decreasing the proton flux. The curvature of the target will also result in a larger overall divergence of the beam compared to that emerging from a flat foil.

Even though the pulsed high-field solenoid lenses were demonstrated to have better acceptance and transmission compared to quadrupole lenses, they still suffer from limitations when operation is required at higher magnetic fields or higher repetition rates. These limitations may be improved with further developments to improve the handling of the mechanical stresses induced by the interaction between the current-carrying wire and the magnetic field it generates, the Ohmic heating at high repetition rates, and the voltage breakdowns.

1.3.3 Plasma focusing

The use of a plasma to focus proton and positive ion beams was proposed several decades ago based on two different devices. In 1947, D. Gabor proposed a device that confines a non-neutral electron plasma to form
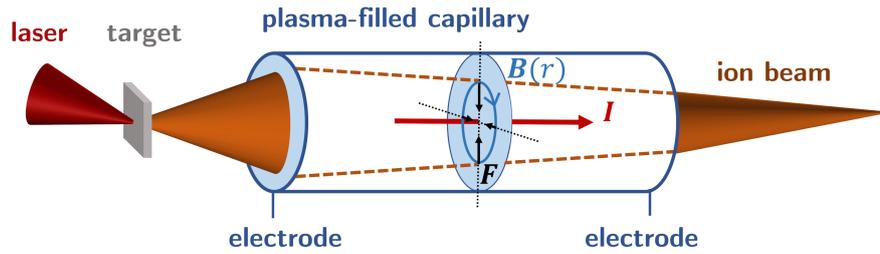
Figure 1.6: Schematic representation of the focusing effect of an active plasma lens (APL) on a positive ion beam. The axial current I is established between two electrodes placed at the ends of a gas-filled capillary. During the plasma discharge, the positive ions passing through the lens are subject to a radially inward focusing force F.

a space-charge lens [49]. The electron plasma provides a radially inward electrostatic force which increases linearly with the radius if the plasma density is uniform, resulting in an ideal beam-focusing effect. In 1950, a different device was constructed to focus $350\,\text{MeV}$ proton beams at BNL [129] with the use of a plasma discharge. A high current was passed through a cylindrical glass tube filled with hydrogen, the axis of which was parallel to the beam. If the current density is uniform, a uniform azimuthal magnetic field that increases linearly with the radius is created and the beam experiences a linear radial restoring force. As the next chapter provides a more in-depth overview of non-neutral electron plasma lenses, a discussion is included here on the suitability of current-based plasma lenses (also called active plasma lenses, APL) for focusing laser-driven proton and ion beams.

In an APL, the proton or ion beam propagates collinearly with an externally driven current I. The radial focusing of the beam is a result of the radial magnetic field gradient g, which for a current with uniform density is

$$g = \frac{\partial B_\theta}{\partial r} = \frac{\mu_0 I}{2\pi R^2},$$ (1.5)

where B_θ is the azimuthal magnetic field strength, r is the radial distance from the axis of the channel, μ_0 is the vacuum permeability and R is the radius of the cylindrical gas tube. The discharge in the lens, and, hence, the peak current, have a duration of around a hundred of nanoseconds which makes the APL suitable for use with a pulsed beam with short bunches.

APLs have been installed on ion beam lines at CERN [37] and GSI [38] to focus different types of ions from MeV Krypton ions to GeV Gold ions. The radii of these APLs are more than a centimetre and, thus, large discharge currents can lead to the magnetic pressure at the tube wall overcoming the plasma pressure, i.e. when

$$\frac{B^2}{2\mu_0} = \frac{g^2 R^2}{2\mu_0} > n_0 k_B T_e,$$ (1.6)

where T_e is the electron temperature, and n_0 the electron density. Under this condition, the plasma contracts radially in the so-called 'z-pinch' effect and tends to concentrate the current near the axis forming a nonlinear

and potentially unstable field that leads to degradation of beam quality.

In recent years, compact APLs based on gas-filled discharge capillaries have been shown to produce high-magnetic gradients with flexible tunability for radially symmetric focusing of electron beams [130–132]. Gas-filled discharge capillaries typically consist of hollow tubes several centimetres long of diameter 250–2000 µm laser machined into a sapphire or glass ceramic substrate. Two electrodes are place on each end to provide the voltage difference and the typical discharge currents are in the range of a few hundred Amps to a few kilo-Amps.

The main challenge for APLs are the different sources of non-linearity in the focusing field, mainly due to nonuniform plasma heating [133] and wakefields [134]. Recent work [132, 135] has shown that linear focusing forces can be achieved and beam emittance preserved by avoiding the plasma nonlinearities. The effect of the wakes may be reduced by decreasing the length of the lens and compensated for by increasing the focusing field gradient. The fields were shown to be linear for focusing gradients in excess of $3500 \, \text{T/m}$ inside a $15 \, \text{mm}$ long capillary with a diameter of $500 \, \text{µm}$ [136].

Unlike electron beams generated from laser-plasma acceleration, laser-driven proton beams are non-relativistic and much more divergent. The relatively small aperture of the APLs based on gas-filled capillaries limits the flux of protons or heavier ions that can be captured from the target. Currently, the development of APLs is an active field of research and capillaries with a diameter of $2 \, \text{mm}$ or a maximum length of $40 \, \text{cm}$ have been successfully tested [39]. Larger diameters of the capillary are believed to be possible with improved vacuum pumping systems [39]. An analytical evaluation [137] of some of the constraints of proton propagation in plasma channels suggests that the main beam characteristics, such as emittance and energy, are not significantly affected by the interaction with the plasma. However, experimental investigations are needed to confirm that linear focusing can be achieved for proton and ion beam transport.

As a first demonstration, an APL was employed to collect and transport protons generated from the interaction of a PW laser with a tape target as a part of an experimental platform, BELLA [25], that can deliver laser-driven proton pulses to biological samples. The $1 \, \text{mm}$-diameter Argon gas-filled capillary has a length of $33 \, \text{mm}$ and produces a focusing gradient up to $600 \, \text{T/m}$. The transport efficiency achieved by the APL at BELLA was $\sim 0.2\%$ for protons of $>1.5 \, \text{MeV}$. Given the $1 \, \text{mm}$ diameter aperture of the APL at $13 \, \text{mm}$ behind the target, the precise alignment of the lens relative to the ion source is crucial to maximise the transmitted number of protons.

1.3.4 Energy dependant focusing

The dynamical behaviour of charged particles in the particular beam-line components discussed above is of interest to establish the focusing capabilities of these systems. Writing down the equations of motion for a charged particle moving through a focusing field allows the dependence of the focusing strength on the key parameters of the beam-transport system to be extracted. For a detailed exposition of the single-particle linear dynamics for a selection of beam-focusing systems, the interested reader is referred to appendix B. Furthermore, the 'chromatic' dependence of the focusing strength on the energy or momentum of the particle is also obtained

Table 1.4: Overview of the focusing strength parameter, k, of a selection of devices that can act as beam-focusing lenses for laser driven ions of mass m, atomic number Z, momentum P_0, and relativistic Lorentz factor γ_0. For the definition of k and detailed derivations see appendix B.

Focusing lens type	Focusing strength[†], k	Lens parameters
Quadrupole doublet[‡]	$e^2 \left(\dfrac{Z^2}{P_0^2} \right) \dfrac{B_0^2}{a^2} ls$	magnetic flux density (at poles), B_0 pole radius (max. beam radius), a
Solenoid magnet	$\dfrac{e^2}{4} \left(\dfrac{Z^2}{P_0^2} \right) B_0^2$	magnetic flux density, B_0
Non-neutral electron plasma lens	$\dfrac{e^2}{2\epsilon_0} \left(\dfrac{Zm\gamma_0}{P_0^2} \right) n_e$	electron plasma density, n_e
Active plasma lens (APL)	$\dfrac{\mu_0 e}{2\pi} \left(\dfrac{Z}{P_0} \right) \dfrac{I}{R^2}$	discharge current, I lens radius, R

[†]Under the thin-lens approximation, the focusing strength, k, is related to the focal length, f, via $1/f = kL$, where L is the effective length of the lens.
[‡]A focusing and defocusing pair of thin quadrupole lenses of length l, separated by a short drift space of length s (from centre to centre) [127].

from the equations of motion. Table 1.4 specifies how the focusing strength depends on the parameters of the specific focusing element and of the beam, for a selection of focusing systems. For example, for a non-relativistic beam, the focusing strength of both solenoid magnets and non-neutral plasma (Gabor) lenses is inversely proportional to the kinetic energy of the particles in the beam. Conversely, the focusing strength of a quadrupole or an active plasma lens is inversely proportional to the square of the beam energy. Consequently, ions with different kinetic energies will experience stronger chromatic effects if focused by solenoids or Gabor lenses compared to the case of using quadrupoles or active plasma lenses. The chromatic effects can be leveraged to filter the protons and ions in a beam based on their atomic number, kinetic energy or momentum.

The chromatic dependence of the focusing parameter k on the energy of a particle that traverses the focusing field is of interest in view of designing a scheme to filter the energy of a beam captured from a laser-driven source. Energy selection is necessary for laser-generated proton and ion beams since they are produced at the source with a continuous energy spectrum which typically has an energy spread $>100\%$. One possible solution, which is also proposed for LhARA [9], is to employ the dependence of the focusing strength of the capture system on the beam energy. Placing a collimator at the focal point of the beam that has the nominal energy removes particles with large energy deviations.

Figure 1.7 shows the strong energy dependence of the focusing strength parameter, k, of a solenoid and a non-neutral plasma lens in comparison to the weaker variation that is characteristic of an active plasma lens. It is important to note that for both protons and C^{6+} ions at low energy, the relative change in the focusing strength is almost identical for a solenoid and a Gabor-type plasma lens. This is explained by table 1.4 (or, equivalently, eqs. (B.13) and (B.17)) which show that, in the non-relativistic regime, the variation of the focusing strength with P_0^2, or energy, is identical. Figure 1.7 shows how the focusing parameter, k, changes as a function of energy

relative to its value at the nominal energy in the case of protons and carbon ions focused by a solenoid, a Gabor lens, and an APL. The asymmetry present in fig. 1.7 shows that particles with a lower energy than the nominal value suffer from a larger variation in the focusing strength. Therefore, the low energy part of the beam can be filtered more easily than the high energy part with a collimator placed at the focal point of the particles with the nominal energy. Due to the difference in the corresponding linear transfer matrices and the additional non-linear effects, a better comparison of the energy-selection capabilities of solenoids and Gabor lenses should be based on particle-tracking with field maps.

1.3.5 Discussion

In conclusion, several systems have been described which are used routinely to capture, transport, and focus laser-generated proton beams. Each of the techniques presented above suffers from at least one of the following limitations: low beam-transmission efficiency, limited repetition rate, or lack of the capability for fast tuning. As part of a future facility that can deliver laser-generated proton and ion beams for radiobiological studies or radiotherapy with a wide range of beam characteristics, the ideal capture and focusing system should overcome these limitations.

In principle, the high beam-loss in the capture section could be mitigated by generating a larger particle flux at the source within the energy interval of interest. Larger particle fluxes can be obtained with higher power lasers that increase the cost of the facility and may not be commercially available. Similarly, conventional superconducting solenoids can provide the required focusing strength required by protons beams of a few tens of MeV. However, a superconducting coil needs further consideration to confirm that magnet quenching can be prevented. The protection against quenching needs to be evaluated due to the close proximity of the laser-target interaction region to the superconducting coil. the large divergence of the particles emerging from the target, and the high intensity of the electromagnetic fields involved in the acceleration mechanism. Furthermore, the superconducting solenoid is embedded in a cryogenic system which would increase the complexity of operation and the overall cost of a research or treatment facility. Considering the limitations described above for the main techniques that can capture and focus laser-generated proton and ion beams, the plasma lens remains a suitable system which could overcome those limitations and meet the demands of a facility for radiobiological studies or radiotherapy.

Two particular areas where laser-driven sources have a disruptive potential to bring ground-breaking contributions are radiobiology research and, eventually, clinical radiotherapy using charged particles. A concept facility that is dedicated to radiobiological studies is LhARA, which has been designed to employ several plasma lenses to make full use of the unique characteristics of the proton and ion beams generated by a laser-driven source. The fundamentals of radiotherapy, along with the concept of the LhARA facility, are described in the next section, with emphasis on the role of the plasma lenses in the delivery of beams with novel time, energy, and spatial structure.
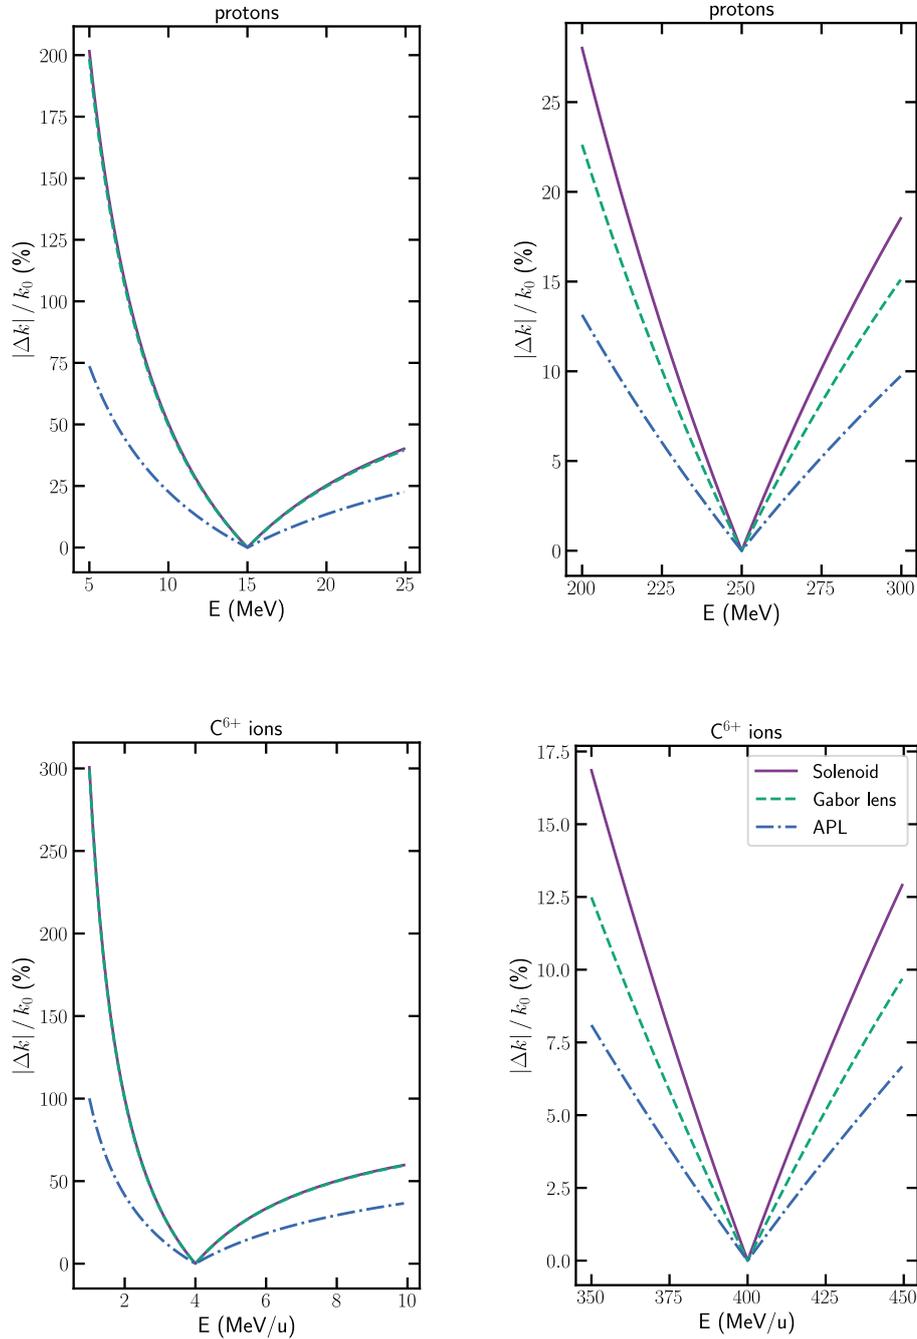
Figure 1.7: The chromatic change $|\Delta k| = |k - k_0|$ of the focusing strength parameter k as a function of the beam kinetic energy for protons and C^{6+} ions. The magnetic field of the solenoid, the electron density of the Gabor lens, and the gradient of the APL are fixed such that at the nominal energy where $\Delta k = 0$ all three focusing elements provide the same k_0.

Chapter 2

Particle radiotherapy and radiobiology research

Radiation therapy (RT) is an advancing treatment modality for cancer and the outcomes of RT have seen substantial improvements in recent years. An important role in this context has been played by charged particle beams, in particular protons and heavier ions, which are widely recognised to produce clinical benefits not achievable with conventional radiotherapy using photons. However, systematic studies are needed to improve the understanding of how the physical characteristics of charged-particle beams generate the biological response of healthy and tumour tissues to ionising radiation.

This chapter describes the fundamental concepts of radiotherapy using charged particles (sections 2.1.1 and 2.1.2), presents the current types of particle accelerators used in treatment facilities (section 2.1.3), and discusses the need for charged-particle beams with novel characteristics (section 2.1.4). Then, an outline of the LhARA facility concept is given (section 2.2.1) with emphasis on the use of plasma lenses envisioned for LhARA to make full use of the unique properties of ion beams generated by a laser-driven source (sections 2.2.2).

2.1 Charged particle therapy

2.1.1 Fundamentals of radiotherapy

The therapeutic window

Ionising radiation can induce damage in both normal and cancerous cells. Nonetheless, radiotherapy exploits the fact that normal cells are generally better at repairing this damage compared to cancerous cells. Thus, the safe and effective delivery of RT is based on finding the optimal balance between the probability for tumour control (TC) and the probability for complications of normal tissues (NTC) [139]. Figure 2.1 illustrates the the probabilities of TC and NTC as functions of the absorbed dose. The response to radiation changes slowly at low or high applied dose, but there is a sharp increase at intermediate dose values. The probability for NTC starts to significantly increase at a higher dose threshold compared to the one of TC. As a result, there is a range of doses for which the probability of TC is greater than the probability for NTC. This interval is called the *therapeutic window*. The horizontal extent of the therapeutic window corresponds to the doses that can be therapeutically
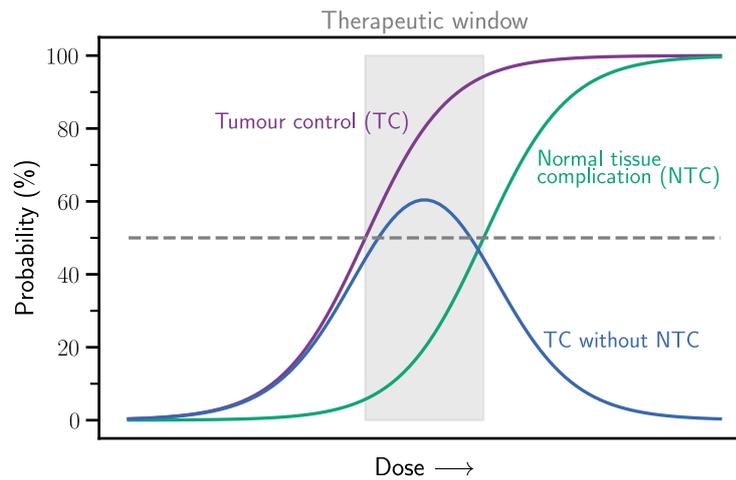
Figure 2.1: Schematic illustration of the probability for tumour control (TC) and normal tissue complication (NTC) as functions of the absorbed dose. The interval of absorbed dose where the probability of TC is high and the probability of NTC is low defines the *therapeutic window* of the doses that can lead to clinical benefits [138]. The blue curve (TC without NTC) indicates the difference between the probability for TC and the probability for NTC as a function of the radiation dose.

exploited, while the vertical size indicates how likely it is for a treatment to be safe for the patient [138]. In the clinical context, there are a variety of approaches to widen the therapeutic window and, thus, to improve the RT treatment, including fractionation schemes (changes in the temporal or spatial distribution of the dose or dose fractions), concurrent chemotherapy, modulation of the immune response, and the administration of drugs acting as radio-protectors or radio-sensitisers [140].

Depth-dose distribution

Ionising radiation deposits energy within the irradiated object as a consequence of the physical interactions of the ionising particles with matter. Photons and neutrons are uncharged particles so they transfer their energy indirectly by liberating charged particles from the atoms within the irradiated object through the Compton effect or nuclear scattering, for example. Protons and other ion species ionise the atoms directly via Coulomb interactions. These charged particles undergo frequent inelastic scattering with the orbital electrons and, thus, are slowed down by a quasi-continuous energy loss up to the point at which the beam completely stops. As a consequence of the different interaction processes through which they interact with matter, each species of ionising particle has a characteristic profile of energy deposition along the trajectory within the irradiated object. The resulting depth-dose profile has strong implications for optimising the treatment planning with a particular type of ionising radiation.

Figure 2.2 shows the depth-dose distributions of several types of radiation used for cancer therapy. For protons and heavier ions, the depth-dose curve exhibits a pronounced maximum towards the end of their penetration depth as the particles come to rest. This characteristic maximum is called the 'Bragg peak' and is typically composed of a steep slope at the proximal edge and a sharp drop-off to zero at the distal edge. The position of
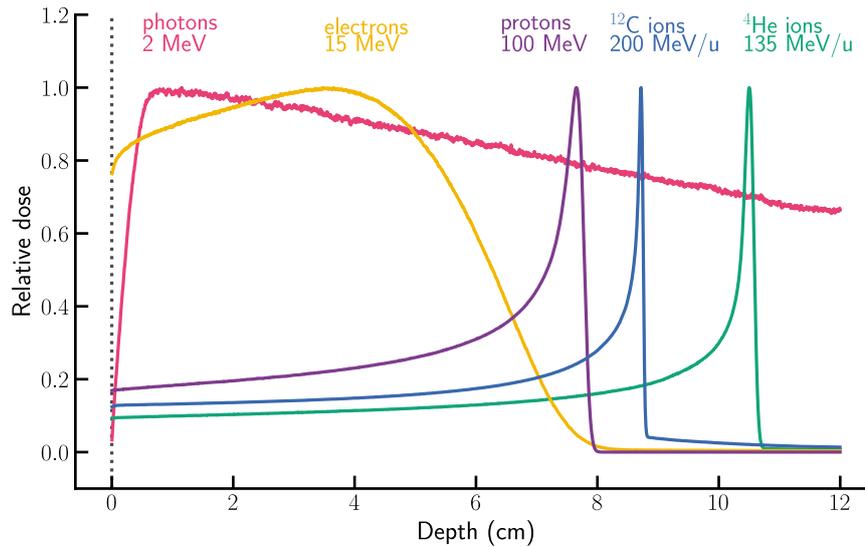
Figure 2.2: Depth-dose distributions (normalised to their maximum) for heavy charged particles such as protons and carbon ions compared to other therapeutically relevant particle types, electrons and photons. The absorbed dose profiles were obtained from Monte Carlo simulations [141] of pencil beams that travel through a water phantom with the beam energy at the entrance indicate in the figure.

the Bragg peak is dictated by the beam energy and its width depends in particular on the spread of energy of the particles in the beam.

The uncharged radiation species (photons and neutrons) exhibit a buildup of the dose deposited near the entrance to the tissue where the dose quickly grows from a low entrance value to its maximum before starting to decrease at greater depths. The dose buildup is a result of electrons being liberated near the surface and depositing their energy as they travel short distances into the irradiated object before stopping. As the beam propagates deeper, particle scattering and absorption cause the dose to decrease monotonously. This long tail in the dose distribution represents an important disadvantage of uncharged particles as it generally results in non-negligible dose being delivered to non-targeted tissues and organs. Nonetheless, the low entrance dose can be beneficial for skin sparing.

Compared to the conventionally used photons, charged particles exhibit a more localised dose distribution which allows a more conformal irradiation of the tumour and sparing of the healthy tissue. This is a consequence of charged particles losing energy at a larger rate as they slow down through the irradiated object. Thus, protons and heavier ions exhibit a relatively low entrance dose that stays nearly constant before the Bragg peak. As the beam slows down, the rate of energy loss increases causing the particles to stop at a certain depth and the dose to drop rapidly to zero. Heavier ions yield a sharper Bragg peak compared to protons and a smaller entrance dose normalised to the peak value. However, non-elastic nuclear interactions of the heavier ions create lighter nuclear fragments that can propagate deeper into the medium. This gives rise to a low-dose 'fragmentation tail' beyond the Bragg peak. The presence of the fragmentation tail can cause severe toxicities in normal tissue for very heavy ions ($Z \geq 12$) which is why the use of ions heavier than carbon or oxygen is not considered anymore

for clinical treatments.

Relative biological effectiveness (RBE)

The level of damage induced by ionising radiation in cells and tissues is the result of a complex mechanism of physical, chemical, and biological processes. Ionising radiation primarily causes breakage of chemical bonds and alteration of molecules through ionisation and excitation of atoms. The most crucial targeted molecule is the DNA of the cancerous cells. The DNA molecule can either be directly affected by the ionising radiation or indirectly damaged by other cellular atoms or molecules that result, for example, from the interaction of the radiation with the water inside the cell. Hence, the processes that dictate the fate of an irradiated cell are directly related to the characteristics of the interaction of a specific particle with matter. Different types of radiation (e.g. X-rays, proton or carbon beams) yield different cellular responses. Thus, for the same absorbed dose, irradiation with proton and ion beams are empirically found to be biologically more effective.

As a consequence, the efficacy of protons, carbon ions and other light ion species is characterised by their relative biological effectiveness (RBE) in comparison to reference photon beams. The RBE is defined as the ratio of the dose delivered by a reference X-ray radiation, $D_{\mathrm{ph,ref}}$, to the dose of the radiation under consideration, D_{ion}, such that both doses result in an identical biological effect for a specific endpoint *X*:

$$RBE\,(\mathrm{endpoint}\,X) = \frac{D_{\mathrm{ph,ref}}\,(\mathrm{endpoint}\,X)}{D_{\mathrm{ion}}\,(\mathrm{endpoint}\,X)}. \tag{2.1}$$

For example, in proton therapy, a generic value of 1.1 is used for the RBE during treatment planning [142]. This means that patients receive a lower dose of protons by about 10% to produce the same therapeutic effect that would be obtained using X-rays.

In addition to the radiation type and the biological endpoint under consideration, the RBE has been shown to depend on physical factors, including particle energy, absolute dose, dose rate and fractionation, but also on biological factors, such as tissue type and the level of oxygenation [143, 144]. For example, *in-vitro* studies with protons showed evidence for a variation of the RBE from around 1 to values over 3 for clonogenic cell survival [145]. Understanding the variations of the RBE and how to incorporate them into the treatment planning is crucial to treat tumours effectively without significantly damaging the healthy tissues and organs at risk. Similarly large variations of the RBE were reported for in-vivo studies with carbon ions with values ranging from about 1 to as much as 5 [146].

2.1.2 Proton and ion therapy

The use of charged particles such as protons and heavier ion species was initially proposed by R. R. Wilson in 1946 [147] laying the foundations of proton therapy. It was previously known that protons and other light charged ions deposit most of their energy at the very end of their range, in the so-called Bragg peak. In contrast to protons and ions, photons, for example, exhibit a slowly decreasing depth-dose profile all along their propagation path.

As a result, a 'low-dose bath' of radiation is deposited in healthy tissue. It was this property of the interaction of protons and ions with matter that R. R. Wilson identified as a potential means to achieve a high dose conformity to the targeted tumour and less damage to the healthy tissue.

In addition to their capability to deliver a much more targeted dose, protons and ions have been associated with further radiobiological advantages as a result of more recent experimental investigations [148]. For example, protons and ions produce an increased number of direct damages to cells and more complex DNA lesions, inhibit the formation of new blood vessels in tumours, thus, preventing their progression, and can lead to a reduction of inflammatory responses [149]. It must be noted that the advancement of dose delivery with X-rays has led to techniques, such as intensity-modulated RT, which can achieve levels of dose conformity that are comparable to the ones obtained with protons [150]. However, such techniques require a high number of X-ray fields delivered at many different directions or even a continuous modulation of the beam intensity, size and angle.

In the last few decades, proton and carbon ion beams have been used predominantly in hadron therapy for treating cancer patients. Thus the development of the accelerator technologies [151] that can provide clinically relevant beams and the improvement of the treatment modalities that make use of them has mainly focused on these two types of particle. Other ion species, such as helium, oxygen, silicon, and neon have been investigated [152–155], but the fragmentation tail, which is characteristic to very heavy ions, was found to cause severe toxicities in healthy tissue. At present, only carbon is used at medical facilities [17] mostly as a prescription for very radioresistant tumours. With an RBE between protons and carbon ions, and a less-pronounced fragmentation tail compared to that of carbon, helium ions (^{4}He) have recently gained a lot of attention for use in radiation therapy [156, 157].

2.1.3 Treatment facility accelerators

The main challenge in providing proton and ion RT to a larger fraction of the world population is to reduce the cost, size and the complexity of operation of the treatment facilities. The key element of a medical facility is the particle accelerator. Most often, it dictates the requirements of the beam delivery system as well as the performance of the treatment facility. Further development is required to ensure that the treatment plan for patients will be determined by clinical need and will not be limited by the available accelerator technology.

In the current state of the field, protons are accelerated to therapeutic energies from 70 to 250 MeV mainly with cyclotrons, while only synchrotrons are used for carbon ions up to 400 MeV/u. Each of these two types of circular accelerators has its advantages and disadvantages. State of the art intensity-modulated proton therapy may require up to 60 different energies, each suitable for dose to be deposited in a single transverse layer of the target volume [158]. Both accelerator technologies above are able to provide a variety of particle energies. However, the beam energy can be changed within an order of a couple of seconds which is inefficient and significantly increases the time that the patient needs to stay still. Furthermore, the use of heavier particles such as helium and carbon requires an increase in the overall size of the accelerator or in the strength of the bending magnets which typically leads to a drastic increase in the cost of the accelerator.

Cyclotrons

The cyclotron was invented in the early 1930s as the first cyclic accelerator [159]. However, it is still one of the most used technologies for accelerator applications today, including proton therapy, due to its great reliability and relatively small footprint. Out of the current 127 operational hadron therapy facilities, 116 use a type of cyclotron [17].

Cyclotrons produce a continuous stream of protons accelerated to the maximum design energy. They are usually more compact and able to achieve high beam intensity. The protons follow a spiral orbit in a plane between the poles of a single magnet ('Dees') while an oscillating RF electric field accelerates the beam between the Dees for a few hundred turns. During the acceleration, there is a significant change in the revolution frequency of protons or ions which is compensated by either increasing the B-field with radius (isochronous cyclotron) or by changing the frequency of the accelerating electric field (synchrocyclotron).

A typical normal conducting cyclotron has a mass around $250\,\mathrm{t}$ and a diameter of about $5\,\mathrm{m}$. A more recent synchrocyclotron from Mevion [160] uses superconducting magnets that reduce the total mass to only $20\,\mathrm{t}$ [161]. Commercial companies are currently pushing the magnetic field up by factors of 2-4 by switching from normal-conduction to super-conducting magnets, which has led to a significant drop in price. Larger reduction in cost and size is believed to be possible in the future by developments that will replace the conventional accelerators based on resistive magnets with superconducting cyclotrons and, reaching eventually to iron-less facilities [162]. Work has started on designs of cyclotrons with no iron yoke, thus substantially reducing the weight and the time taken to change the magnetic field. Both roles of the yoke to shape the magnetic field in the accelerating gap and shield the rest of the cyclotron from the field would be taken over by a system of SC coils. However, cyclotrons with strong magnetic fields can only operate in pulsed mode (synchrocyclotrons). The time structure of such a pulsed beam imposes minor limitations to the spot-scanning technique, but it does not allow for the more advanced continuous line scanning.

The main disadvantage for proton therapy is that cyclotrons cannot provide the fast beam energy modulation required for the beam-scanning technique. This is due to the use of energy degraders which are mechanically inserted in the path of the beam before the treatment room. This method further requires an energy selection system and increased radiation shielding. To compensate for the particle loss inside the energy degraders, a higher circulating current is combined with a good extraction efficiency. Particle loss of up to 99% is possible when degrading to the lowest energy transported and collimating the beam for proper beam-size and divergence. Thus, the cyclotron becomes radioactive during operation due to beam losses, especially for a low extraction efficiency, and, thus, requires more complex radioprotection systems.

A further major disadvantage of cyclotrons is that energies above $250\,\mathrm{MeV}$ for protons, which are desirable for imaging the patient during the treatment, are not achievable while maintaining a compact size. Also, while significant efforts have been made to design a cyclotron suitable for carbon therapy [163], the plans for treatment facilities based on this particular accelerator technology have stalled over the past decade.

Synchrotrons

Synchrotrons accelerate pulses of particles to the desired energy. The particle bunches are initially accelerated by a linac and then injected into the synchrotron. Typically, the ring is made of four to eight dipoles with quadrupole and sextupole magnets in between them, as well as an RF cavity. The magnetic field and RF-frequency are synchronised with the revolution frequency of the particles such that the radius of the orbit remains constant during acceleration. Once the specified energy is achieved, a particle bunch is extracted from the synchrotron into a transport line towards the treatment rooms. The beam produced is characterised by a dead time of 1–$2\,\mathrm{s}$. This is required to ramp down the magnetic field and re-accelerate the next particle bunch. Thus, the energy can be adjusted cycle by cycle. In general, for the application of one treatment layer of radiation at the patient, one to three filling and acceleration sequences are required.

The main advantage of synchrotrons is the ability to provide active energy modulation and smaller energy spread with lower power consumption [18]. Since the path of the particles is fixed, synchrotrons are built with much smaller magnets than cyclotrons, even though additional quadrupoles are needed for the strong focusing. The extraction may occur on a relatively long period of time (up to 5 s) during which the magnetic field is kept constant (flat-top). This allows the beam to be scanned at the patient in a similar manner to a continuous beam. Energy changes reach a time scale of milliseconds as the energy of a slow extraction synchrotron can be varied during the extraction of a single pulse. Another advantage is that because there is no need for an energy degrader, the average current required is greatly reduced. Furthermore, synchrotrons can accelerate particles with larger magnetic rigidities. Thus, they are the only technology used at present in operational facilities for carbon therapy such as HIT, CNAO, MedAustron and NIRS [164].

An indirect measure of the beam size is given by the beam emittance [125]. While the emittance obtained with a synchrotron is well suited for treatment, space charge effects limit the maximum achievable beam current. Storing more charge is possible if the injection energy is increased. However, this drives the cost up since a longer pre-accelerator is needed. Another challenge for a facility which uses a synchrotron is the treatment of targets inside moving organs. The main strategy is based on the ability to start and stop the beam quickly enough to deposit the charge in the right location. Some fast extraction schemes [165] have been proposed which can achieve repetition rates as high as $30\,\mathrm{Hz}$ with the beam fully extracted in one turn. However, to deliver the dose with 1% accuracy, the energy change and beam scanning needs to be done in about $20\,\mathrm{ms}$ [164], still faster than currently possible.

Multiple optimisation techniques have been developed to build smaller and cheaper synchrotrons. Better magnet designs with an increased ramping speed allow for the decrease in the dead time between spills. The power consumption has been reduced by improving the injection chain using a Radio Frequency Quadrupole (RFQ) directly injecting into the synchrotron. Modifications of the beam optics and field-regulation of the magnets [166] further improved the treatment time. Only recently, superconducting magnets have been investigated for use in medical synchrotrons at NIRS (Japan) where a reduction in diameter from 25 m to 7 m is possible

for a carbon-ion accelerator. At HIT (Germany), the efforts have been focused on the stability of the extracted beam intensity which is crucial for continuous scanning. However, no substantial size reduction is expected in the near future for proton therapy facilities driven by a synchrotron. The cost reduction is possible with further optimisation of magnets, RF cavities or injection schemes, but the process contains many small slow steps.

2.1.4   Novel beams for radiobiology

Particle beam therapy (PBT) is routinely delivered today at low dose rates ($<10\,\mathrm{Gy/min}$) with beam characteristics restricted to a small number of temporal schemes, and a small number of spatial distributions. A conventional treatment plan consists of 'fractions' of $\sim 2\,\mathrm{Gy/day}$ delivered over several weeks in daily sessions. The dose is delivered predominantly with large homogeneous beams that irradiate several square centimetres. In one such treatment, the damage to the healthy tissue limits the dose delivered and, thus, constrains the clinical efficacy of PBT. The limited tolerance of the healthy tissue leads to treatments with limited effectiveness especially for radioresistant tumours and paediatric cancers. The use of novel beams with strikingly different characteristics has led to evidence of enhanced therapeutic benefit. However, the radiobiological mechanisms that maintain the therapeutic effect while significantly reducing damage to healthy tissue still have to be identified and explained. Studying these new regimes of PBT and testing them systematically in treating patients provide a path to a radical transformation of PBT.

FLASH radiotherapy

A potentially highly beneficial differential response of tumour and healthy tissue to ionising radiation was reported in the so-called 'FLASH' regime [12]. In FLASH radiotherapy, the dose is delivered in a very short amount of time (of the order of tens of milliseconds as opposed to several minutes in conventional RT). As the total dose is comparable to conventional RT (1–$10\,\mathrm{Gy/pulse}$), the average dose rate in FLASH RT is exceedingly high: $\geq 40\,\mathrm{Gy/s}$. Irradiating animals in this way has proven to yield equivalent tumour control to standard RT while significantly reducing the side effects in normal tissues [167, 168], a phenomenon known as 'FLASH effect'. Moreover, FLASH is advantageous due to the short irradiation time which bypasses problems related to internal organ motion during the treatment and results in shorter overall treatment time. The biological advantages of FLASH have been confirmed in mini-pig models and cat cancer patients [11]. More recently, the first human patient was successfully treated with FLASH RT using an electron beam [169]. Furthermore, the first in-human-use of proton FLASH RT was very recently reported as part of a clinical trial of ten patients with bone metastases [170]. This preliminary proof of the clinical feasibility of very high dose-rate protons supports the further exploration of this regime.

The underlying mechanism responsible for the FLASH effect is still not fully understood, but the most current hypothesis is based on the rapid depletion of oxygen at high dose rates which creates a transient radiation-induced hypoxic environment [171]. The local absence of molecular oxygen decreases the degree of non-reparable DNA damage, and thus reduces radiosensitivity. The recent interest in understanding the underlying mechanism is yet

to be matched by the availability of particle beams in the FLASH regime. In addition to the electron beams that have been predominantly used so far [13] and which have low tissue penetration, comprehensive *in-vivo* studies are needed with high-dose-rate proton beams in order to translate FLASH RT into clinical practice.

As an integrated part of the research into the FLASH effect, ultra-high dose rate radiobiology is a relatively new field that could be explored using the unique characteristics of laser-generated proton beams, such as the ultra-short pulse duration and the high number of protons per pulse. Recent experimental investigations have reported the delivery of laser-driven proton pulses with dose rates between 10^7 and $10^9 \mathrm{Gy/s}$ for single-shot irradiation which are significantly higher than the conventional ones in the context of FLASH (between 40 and $1000\,\mathrm{Gy/s}$) [23]. However, a continuation of the research is needed to establish whether the differential biological response observed in the FLASH regime extends to laser-driven protons and ions delivering ultra-high dose rates.

Spatially fractionated radiation therapy

In addition to the new temporal schemes of beam delivery that generate the FLASH effect, therapeutic benefit has been demonstrated through the use of heterogeneous spatial distributions of the delivered dose. In standard radiotherapy, the tissue is irradiated homogeneously with the use of large beam sizes ($>1\,\mathrm{cm}^2$). A new RT paradigm proposes strong spatial modulation of the dose deposition in a therapeutic strategy called 'spatially fractionated RT' (SFRT). Based on the sizes of the narrow beams and their spatial arrangement, SFRT includes several techniques [172], such as micro- and mini-beam radiation therapy, GRID, and lattice therapy. In mini-beam radiation therapy, narrow beams ($<1\,\mathrm{mm}$) result in areas of high dose ('peaks') which are spaced by areas of low dose ('valleys') of typically a few mm.

Spatial fractionation of the dose has been shown to increase significantly dose tolerance and sparing of the normal tissue [14, 173] while maintaining or even improving the tumour control rates when compared to those obtained with conventional homogeneous irradiation [15, 174]. In particular, important increases in the therapeutic index have been achieved using micro- and mini-beam radiation therapy. More recently, the use of protons [175] and other heavier ions [176] for mini-beam RT has generated a lot of interest and led to an increasingly active field of research [177]. Part of this research is the very important goal of completely understanding the radiobiological mechanisms underlying the effects observed with SFRT techniques. The direct or indirect participation of several actors [177] has been shown, but systematic investigations are still needed to completely establish the importance of the beam parameters on the therapeutic benefit. For example, two effects that were reported following SFRT are the activation of the immune system [178] that led to radiation-induced responses in unirradiated tissues away from the target and the fast regeneration of capillary blood vessels in the path of the mini-beams with the help of undamaged cells from the valley regions [179].

Two of the most important indicators of the potential for normal-tissue sparing are the absolute dose delivered within the valleys and the peak-to-valley dose ratio. Therefore, the generation of the narrow beams is a crucial aspect of the micro- and mini-beam RT as it determines the shape and size of the beamlets, and thus, the spatial

distribution of the dose. Based on current experimental data, pencil-shaped mini-beams with sub-millimetre FWHM (<1–$1.5\,\mathrm{mm}$) [177] are believed to maximise the therapeutic index. Current proton therapy facilities are not designed to deliver such small beam sizes, as the current lower limit for pencil beams is $\sim 4\,\mathrm{mm}$ [180]. At clinically relevant energies, proton and ion mini-beams have been generated up to now with the use of multislit collimators which are inherently inefficient as they absorb part of the beam and introduce a source of secondary neutrons close to the patient.

An alternative method has been recently proposed [177] to generate charged particle mini-beams based on magnetic focusing in order to use the entire beam for dose deposition, reduce the treatment time and limit the neutron contamination. With this technique, appropriate mini-beam sizes at the target can be obtained if the incoming beam possesses a small emittance (more precisely a small divergence $<0.1\,\mathrm{mrad}$ or strong linear correlation of the beam distribution in the transverse phase-space). Highly laminar proton and ion beams with emittances significantly smaller than those produced by existing medical accelerator facilities are generated by laser-driven sources. Bridging the gap between the unique beam characteristics of laser-driven ion sources and the capability of generating mini-beams with a magnetic-focusing nozzle requires an efficient beam-transport line. The beam line must be able to capture the particle flux from the source and prepare the beam for injection into the nozzle with low overall beam losses while maintaining the small emittance from the source.

2.2   LhARA

The Laser-hybrid Accelerator for Radiobiological Applications (LhARA) is a planned particle beam facility [9, 10] dedicated to radiobiological research. The LhARA collaboration's long-term vision is to transform the delivery of proton- and ion-beam therapy by harnessing the unique properties of laser-driven ion beams. On a shorter term, the LhARA facility has been conceived as a new, highly flexible source of radiation that can contribute to systematic studies of radiobiology with the exploitation of ultra-high dose rates, multiple ion species, and novel temporal- and spatial-fractionation schemes.

Systematic studies are needed to improve the understanding of how the physical characteristics of particle beams generate the biological response of healthy and tumour tissues to ionising radiation. The efficacy of proton and ion beams is known to depend on many physical factors, including particle energy, dose and dose rate, but also biological factors, such as tissue type or the level of oxygenation. A better understating of the influence of each of these factors can lead to better patient-specific treatment plans that correctly calculate the dose required to treat a particular tumour while minimising the damage to the healthy tissue and organs at risk. LhARA has been designed to allow the study of the potential therapeutic benefit of beams with properties that are not available at conventional facilities.

In addition to the programme of radiobiology, LhARA aims to prove accelerator technologies that can lead to more compact and cost-effective future treatment facilities using charged particle beams. The main novel technologies proposed by LhARA that could, thus, contribute to a scale-up in the provision of proton-
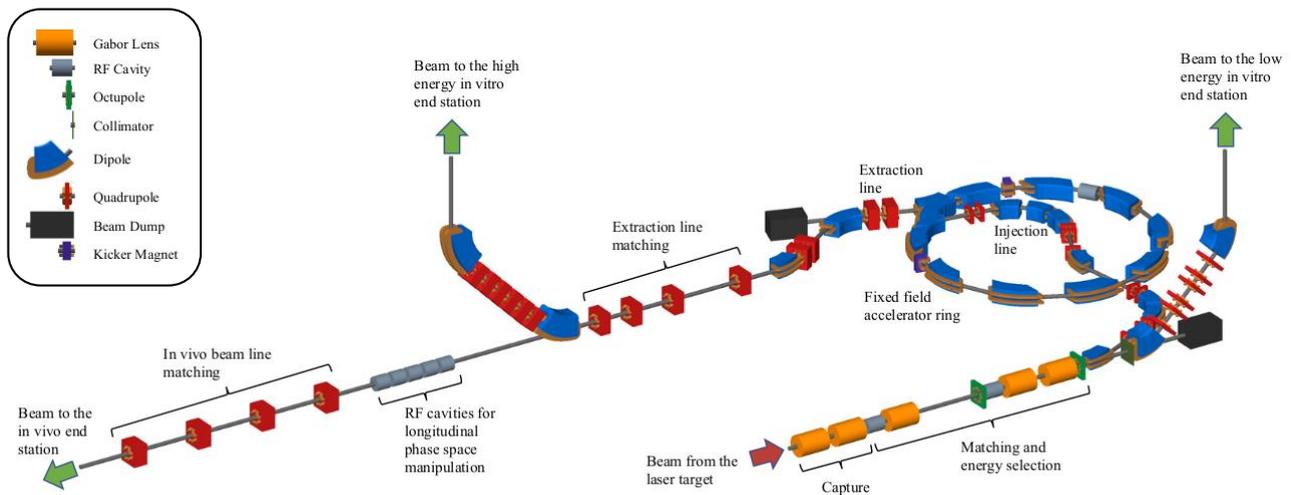
Figure 2.3: Schematic diagram of the LhARA beam lines [9]. The red arrow indicates the particles emerging from the laser-driven source. A series of electron-plasma lenses are used to capture and transport the beam efficiently into a first vertical arc that leads to the low-energy *in-vitro* end station. Post-acceleration is performed using an FFA ring to, then, transport the beam to either a high-energy *in-vitro* or an *in-vivo* end station. The various beam line elements are specified in the legend.

and ion-beam therapy are: a laser-driven proton and ion source, electron-plasma lenses and compact RF post-acceleration.

2.2.1 The LhARA facility

In order to serve a systematic programme of radiobiological studies, the LhARA facility has been designed [9, 181] starting from the concept of a laser-hybrid accelerator. The unique features of a laser-driven proton and light-ion source are planned to be combined with a RF post-acceleration stage to allow both *in-vitro* and *in-vivo* studies. The development of the facility has been conceived in two stages. In Stage 1, the protons up to $15\,\mathrm{MeV}$ from the source are planned to be delivered to a low-energy *in-vitro* end station. In Stage 2, proton and light ion beams are envisioned to be accelerated to higher energies and delivered to one high-energy *in-vitro* and one *in-vivo* end station.

A schematic diagram of LhARA is presented in fig. 2.3 which includes the main elements of the facility and identifies their role in the beam transport. The key components that allow LhARA to deliver beams with a wide range of characteristics are: the laser-driven proton and light ion source based on the TNSA mechanism, electron-plasma lenses for capture, beam focusing and energy selection, and a Fixed-Field Alternating Gradient (FFA) accelerator for compact and fast post-acceleration.

Stage 1

Stage 1 of LhARA consists of all the beam-line components from the laser-driven source to the low-energy *in-vitro* end station and a short abort line. Figure 2.4 shows a diagram of Stage 1 and indicates the relative

position of the main components. Stage 1 has been designed [9] to capture the protons or ions produced by the source, focus them into a beam with low divergence, filter the energy spread of the particles and transport them with low losses through an arc to the vertical end station.

LhARA plans to operate a multi-TW laser that can deliver $>2.5\,\mathrm{J}$ of energy in $<25\,\mathrm{fs}$ pulses at a repetition rate of up to $10\,\mathrm{Hz}$. The laser can drive the acceleration of protons and light ions from the target via the Target Normal Sheath Acceleration mechanism. Due to the need to replenish the target between shots at high repetition rate, a tape target system is envisioned to be developed and optimised [10]. The nominal energy of $15\,\mathrm{MeV}$ for protons was chosen to be lower than the maximum energies that can be achieved by a laser with the parameters listed above. The cut-off energy and, in turn, the flux of protons with energy near the cut-off value, are known to be prone to fluctuations from one laser pulse to another [182]. Scaling down the nominal energy below the achievable cut-off energy is the first step in the development of a strategy to improve the shot-to-shot stability and to obtain a larger particle flux from the source within the nominal energy range.

In total, there are five non-neutral electron plasma lenses that provide the focusing force required to reduce the initial transverse momentum of the particles emerging from the source and to match the beam into the vertical arc. A total of three collimators are located at key positions along the beam-line to ensure a reduction in the energy spread of the initial beam. Two RF cavities are used to manipulate the longitudinal phase space of each particle bunch. The first cavity is placed upstream of the first slit, also referred to as the 'energy collimator', to increase the number of particles within the energy range that is transmitted through the slit. A second RF cavity placed downstream of the energy-selection slit allows the manipulation of the bunch length for variable irradiation time from a pulse at the end station. Furthermore, octupole magnets are used to improve the spatial uniformity of the delivered dose.

The vertical arc consists of two $45°$ bending dipole magnets and six quadrupole magnets. It has been designed to cancel the vertical dispersion and ensure that the first-order transfer map through the arc is equivalent to the identity transformation. Hence, the beam properties at the end of the arc closely match the ones at the entry plane. Varying the characteristics of the beam that is delivered to the end station is possible by changes in the transport properties of the elements upstream of the vertical arc. Finally, the arc is followed by a $2\,\mathrm{m}$ drift tube to penetrate the concrete shielding and bring the beam to bench height. If the first dipole in the arc is not energised, the beam is directed into a beam dump as part of the abort line of Stage 1.

Stage 2

The main role of the Stage 2 of LhARA is to accelerate the beam to energies higher than those afforded by the laser-driven source. The higher-energy particles penetrate deeper inside tissues. Thus, a post-acceleration stage is required to allow *in-vivo* studies on small mammals to be carried out in LhARA in addition to those on cell cultures. For post-acceleration, a fixed-field alternating-gradient accelerator (FFA) based on the spiral scaling principle [183] has been chosen in the baseline design [9] for its many advantages in the context of medical and radiobiological applications. For example, FFAs have the capability to deliver high and variable dose, rapid
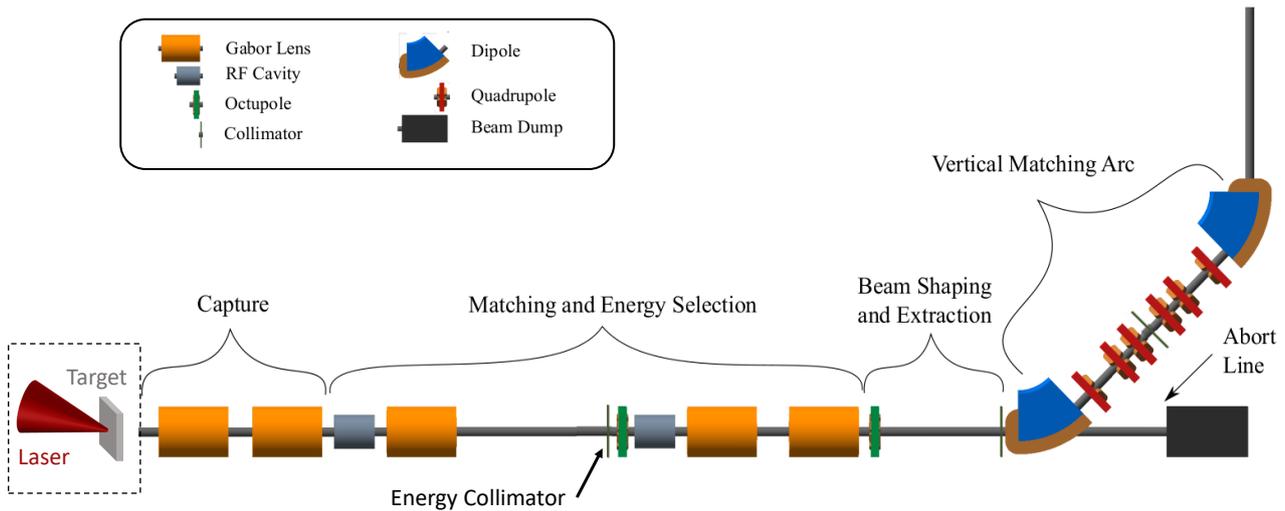
Figure 2.4: Schematic diagram of Stage 1 of LhARA. The beam is created in the laser-target-interaction vacuum chamber just before the capture section of the beam-line. Then, the particles propagate to the right and through the 90° vertical arc to reach the *n-vitro* low-energy end station. The legend specifies the elements of the beam line.

cycling, and the ability to deliver variable beam energies without the use of energy degraders. Furthermore, FFAs can accelerate multiple ion species.

A typical FFA is able to increase the beam momentum by a factor of three, though a greater factor may be achieved (see, for example, [184] or [185]). Thus, for LhARA, the 15 MeV protons from Stage 1 can be accelerated up to a maximum energy of 127 MeV to be delivered to the two end stations of Stage 2. For carbon ions (C^{6+}) with the same beam rigidity, a maximum energy of ~ 33.4 MeV/u can be delivered. The fixed relative boost of the momentum in the FFA implies that the energy at injection into the accelerator determines the beam energy at extraction. In order to change the injection energy, the focusing strengths of the beam-line elements of Stage 1 will be adjusted. For this procedure, varying the strengths of the five plasma lenses is essential as this will allow the required energy slice from the broad spectrum produced at the source to be selected and transported to the FFA with low beam losses. The strengths of the plasma lenses need to be modified from their Stage 1 configuration to produce a nominal beta-function value of 50 m at the entrance to the switching magnet. This scheme simplifies the injection and extraction systems since their geometry and location can be kept fixed.

After acceleration in the FFA, the protons or ions are extracted either towards a 90° vertical arc that delivers the beam to a 'high-energy' *in-vitro* end station or towards the *in-vivo* end station. The transport line for the beam dedicated to high-energy *in-vitro* studies is a scaled version of the transport line to the low-energy *in-vitro* end station in Stage 1. As is noticeable in fig. 2.3, the magnet lengths were scaled up in the high-energy version of the transport-line to keep the magnetic field strengths below the saturation limits of normal conducting magnets [9]. At the end of the straight transport line towards the *in-vivo* end station, four quadrupoles provide the variable focusing required to adjust the beam size delivered at the end station between 1 mm and 30 mm. Further space is reserved for scanning magnets or a magnetic-focusing nozzle [177] that can produce mini-beams.

2.2.2   Role of plasma lenses

A key component of the LhARA beam-lines are the plasma lenses which have been chosen as the baseline solution for focusing protons and ions to address the issues of large divergence and energy spread of the particles accelerated via the TNSA mechanism. The present configuration of lenses [9] plays two main roles in the transport of particles from the source to the end station or to the post-acceleration stage. Firstly, the lenses are envisioned to provide the strong focusing required to capture protons and ions efficiently from the laser-driven source. Secondly, the electrostatic focusing generated by the plasma, in combination with a beam collimator, allows the filtering of particles based on their kinetic energy. More details on how the plasma lenses may achieve these two roles are given below.

Proton and ion capture

The LhARA facility aims to provide a wide range of beam characteristics in both the temporal and spatial domains and to deliver a relatively high number of protons or ions per bunch. In order to achieve that, the facility is designed to make use of the unique properties of the laser-driven source. Hence, the set of key requirements for the beam line of LhARA include the efficient capture of the particles from the source, the formation of a beam with a small energy spread and the transport of the beam to the end stations with low losses. Furthermore, the beam line must be capable of transporting the protons and ions at the repetition rate of the source of up to $10\,\mathrm{Hz}$. Given the requirements above, the baseline design of LhARA employs a series of five non-neutral electron plasma lenses of the Gabor type [49] to capture, focus and transport the beam into the vertical arc of Stage 1 or the FFA injection line of Stage 2. Compared to conventional beam focusing techniques, the use of plasma lenses significantly reduces the mass, size, and cost of the system. At the same time, simulations for clinically relevant energies have shown that the beam energy spread can be reduced by a factor of about three compared to that achievable using conventional systems [66]. Furthermore, the plasma lens provides simultaneous focusing in both transverse planes and it is suitable for operation at a high repetition rate while also providing fast tunability of the focusing strength. The role of the five lenses in LhARA is indicated in fig. 2.4.

As the beam propagates from the solid target, the transverse size of the beam increases rapidly due to the space-charge repulsion of the particles in the beam and the large opening-angle generated at the source. In order to capture a large fraction of the protons and ions emerging from the target, the first plasma lens needs to be placed as close to the source as possible. The separation between the target and the electron plasma inside the first lens is limited by the different vacuum conditions required inside the target vessel and the lens. In the baseline design, a conical nozzle is introduced between the two vacuum vessels to ensure a differential pressure drop. Further space is required after the nozzle and before the entry plane of the plasma within the first lens to fit a cylindrical electrode which is essential for confining the plasma. After the protons and ions are captured by the first plasma lens, a second lens further focuses the particles to generate a parallel beam. Two lenses were chosen to bring the beam from point to parallel in order to decrease the plasma density below the value that

would be required if only one lens was used. Moreover, two lenses provide more flexibility as there is a second focusing strength that can be tuned compared to only one in the case of a single lens within the capture section.

Once the beam is parallel, a third plasma lens focuses the particles back to a small beam size at the location of a beam collimator (termed 'energy collimator' in fig. 2.4). Two additional plasma lenses are employed downstream of the collimator to bring the beam back to parallel and match it into the vertical arc of Stage 1 or the FFA injection line of Stage 2. The two plasma lenses in the matching section can be replaced by four quadrupoles. The quadrupoles only focus the beam in one of the two transverse directions at a time compared to the lenses which focus simultaneously in the two transverse planes. Thus, the quadrupoles can be used to avoid a small beam size in both transverse directions and mitigate potential emittance-growth from space-charge forces. However, the baseline design with the two lenses allows the tuning of the desired beam size at the *in-vitro* end station over a wider range. Only the lower bound on the desired beam size has been achieved so far with the alternative design based on quadrupoles [9].

Beam energy spread reduction

The front end of the Stage 1 beam line has also been designed to filter the protons or ions in the beam depending on their energy and momentum. As a result, the energy spread of the beam at the *in-vitro* end station can be greatly reduced compared to the energy spread of the protons at the source. A combination of electrostatic space-charge focusing (energy selection) and magnetic bending of the beam (momentum selection) allows protons with a narrow energy spread to be delivered to the end station. The main principle of the energy-selection scheme [186] is based on the chromatic dependence of the focusing strength of an electron plasma lens on the momentum and energy of the particles passing through the lens (see eq. (B.17) and fig. 1.7). In the non-relativistic regime, the focal length of an electron plasma lens is proportional to the beam kinetic energy.

In the absence of space-charge forces, the third lens downstream of the source focuses each monoenergetic component of a beam to a single point. With space-charge, the point is enlarged to a finite spot where the beam has a waist. The position of the focal point or the waist of the beam is dictated by the energy of the particles in the beam. A slit placed at the focal point of the particles with the nominal energy absorbs most of the particles with large energy deviation from the nominal value. Figure 2.5 shows the effect of the slit diagrammatically. The goal of the energy selection scheme is to reduce the energy spread from more than 100% at the source to $\leq 2\%$ around the nominal energy which is between 10 and $15\,\text{MeV}$. In addition to the energy collimator placed downstream of the third plasma lens, two other slits are placed at key locations along the Stage 1 beam line. A second circular collimator before the vertical arc shapes the transverse cross-section of the beam to match the circular aperture through which the beam enters the end station. A third elliptical slit is placed in the middle of the vertical arc, in a position with large beam dispersion in one of the transverse directions, in order to remove particles with large momentum deviation that were not absorbed by the energy collimator. The position of the three slits used for energy selection is also indicated in fig. 2.4 relative to the other components of the Stage 1 beam line.
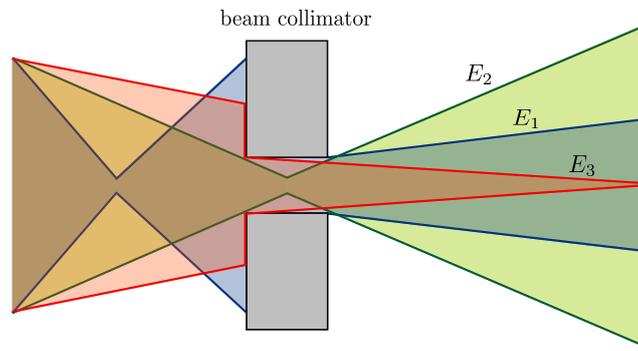
Figure 2.5: Schematic representation of the energy collimation of a beam which has three monoenergetic components $E_1 < E_2 < E_3$. The particles are focused upstream of the collimator by a non-neutral electron plasma lens. The collimator is placed at the focal point of the particles with kinetic energy E_2.

The following chapter presents an overview of the non-neutral electron plasma lens as a device to capture and focus laser-driven proton and ion beams.

Chapter 3

Non-neutral plasma (Gabor) lens

The central subject of this doctoral thesis is the investigation of plasma (Gabor) lenses as focusing devices for proton and ion beams. This chapter introduces the basic principles and general aspects of non-neutral plasma trapping with a focus on the parameters of interest that dictate the focusing capabilities of a plasma lens. First, the fundamentals of non-neutral plasma confinement are presented (sections 3.1.1 and 3.1.2) and the principles of how charged-particle beams are focused by the plasma are explained. Then, a separate section is dedicated to the historical developments in using plasmas to focus charged particle beams (section 3.2.1) and to an overview of the beam-transport capabilities achieved so far (section 3.2.2). The chapter is concluded by a discussion of the attempts to create a practical plasma lens based upon the original Gabor design (section 3.2.3).

3.1 Fundamentals of single-component plasma trapping

Non-neutral electron and positron plasmas can be confined for long periods of time using only static external electric and magnetic fields. The fields are arranged in a particular configuration, shown in fig. 3.1, to ensure that the motion of the particles is restricted to the internal volume of the trap. In an equilibrium state, the confined plasma rotates as a rigid body about the axis of symmetry of the trap. The following two subsections introduce the theory behind the conventional techniques to confine the plasma in the longitudinal direction and in the transverse plane of a Penning-Malmberg trap [188]. A more comprehensive review of the theory of single-component plasmas in electromagnetic traps can be found in [187] and [189].

The ensembles of particles considered below are always assumed to be in the plasma regime which is defined by a plasma length, $L_p > \lambda_D$, and radius, $r_p > \lambda_D$. The Debye length, λ_D, gives the distance over which any external electrostatic potential is screened by the self-organisation of the plasma. A further requirement of the plasma regime is for the particle number density, n, to satisfy the condition $n\lambda_D^3 > 1$.

3.1.1 Longitudinal confinement

The confinement of a non-neutral plasma in the axial direction, z, is achieved by creating an electrostatic potential well. The external electrostatic potential, ϕ_{ext}, is generated by cylindrical electrodes that are individually biased
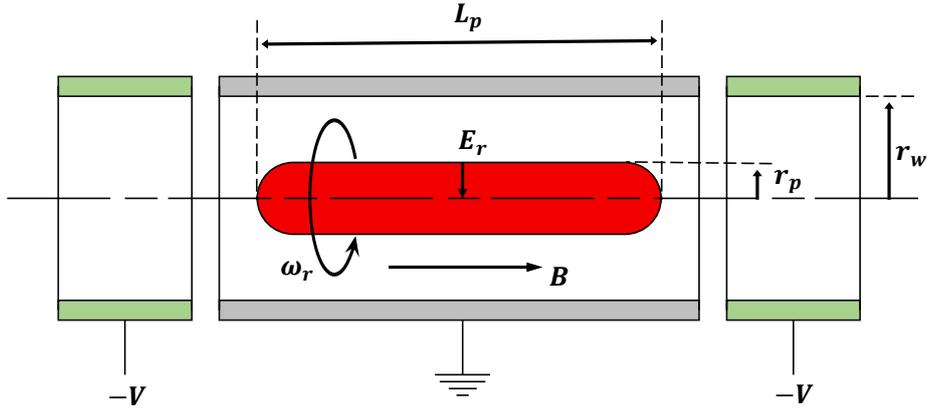
Figure 3.1: Schematic of a cylindrical Penning-Malmberg trap for negative charges (reproduced from [187]). In an equilibrium state, the plasma rotates as a rigid column about the axis of symmetry of the trap as a result of the self-field E_r and the external magnetic field B.

to specific voltages. The electrodes at the two ends of the trap, usually referred as 'end caps', slow down the trapped particles and reflect them back towards the central region of the trap. Due to the Debye shielding, the plasma screens out the axial electric field by its own space-charge potential, ϕ_p, such that the total electrostatic potential, $\phi = \phi_{ext} + \phi_p$, is independent of z within the plasma.

In order to confine the plasma longitudinally, the external potential must be sufficiently large to balance the self-potential of the plasma, ϕ_p. This condition results in a minimum voltage that needs to be applied to the end electrodes to trap the plasma. Consider the case of a plasma column of uniform charge density, long compared to the electrode radius, r_w, and with a rectangular density profile:

$$n(r) = \begin{cases} n = \text{const.}, & \text{for } r < r_p \\ 0, & \text{for } r_p \leq r. \end{cases} \tag{3.1}$$

For an infinitely long plasma column, the electric field inside and outside of the plasma is purely radial and is given by

$$E_r = \begin{cases} \frac{qn}{2\epsilon_0} r, & \text{for } r < r_p \\ \frac{qn}{2\epsilon_0} \frac{r_p^2}{r}, & \text{for } r_p < r, \end{cases} \tag{3.2}$$

where n and r_p are the plasma number density and radius. Assuming $\phi(r_w) = 0$, integration of the radial electric field yields an expression for the potential inside the plasma

$$\phi(r) = \frac{qnr_p^2}{4\epsilon_0} \left[1 + 2\ln\left(\frac{r_w}{r_p}\right) \right] - \frac{qnr^2}{4\epsilon_0}, \qquad r < r_p. \tag{3.3}$$

Equation (3.3) also specifies the potential on axis at the ends of the charge column since the total potential
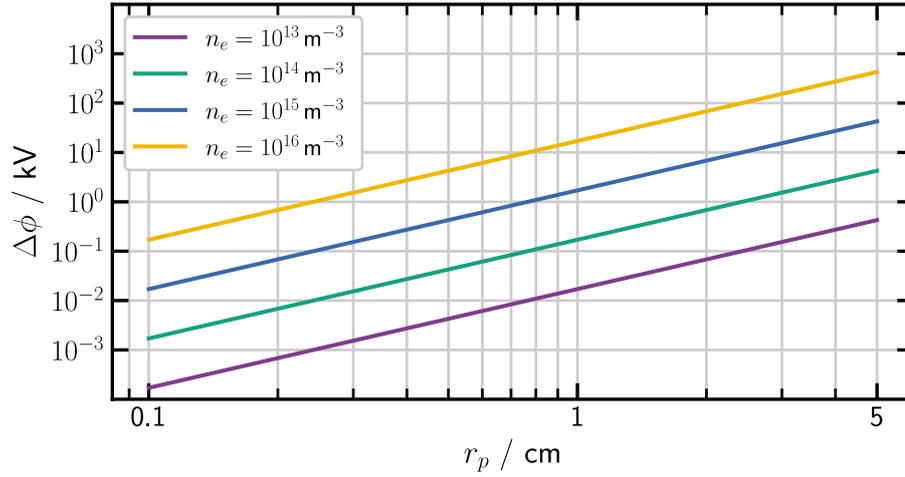
Figure 3.2: The minimum required voltage $\Delta\phi$ that must be applied to the end electrodes of a cylindrical electron trap to successfully confine the plasma longitudinally as a function of the radius, r_p, and the number density, n_e, of the plasma. The radius of the electrodes was chosen as $r_w = 4r_p$.

must be independent of z. Thus, at the ends of the plasma, the potential difference between the wall of the end electrode and the central axis is

$$\Delta\phi = \frac{qnr_p^2}{4\epsilon_0}\left[1 + 2\ln\left(\frac{r_w}{r_p}\right)\right]. \tag{3.4}$$

The voltage applied to the end electrodes of the trap must be at least as large as $\Delta\phi$ to confine the plasma longitudinally. More precisely, the end electrodes must produce a potential well on the axis of the trap with a minimum 'depth' of $\Delta\phi$ for the plasma to be confined axially. For a cylindrical end electrode that is long compared to its radius, a voltage of $\Delta\phi$ is just sufficient for confinement, while for shorter electrodes, a larger applied voltage is required. The aspect ratio of the electrode determines the curvature of the electric field lines and by how much the potential on the axis of the electrode differs from the potential applied to the surface of the electrode. Figure 3.2 shows the minimum voltage required to confine the plasma axially for several electron plasma densities and radii.

For a long, uniform, and cylindrical electron plasma confined with cylindrical electrodes that fully fills the trap at the Brillouin limit (see eq.(3.10)), the space-charge potential on the axis of the plasma, and, thus, the minimum required voltage is given by [187]

$$\Delta\phi = 2.2 \times (B\,[\text{mT}]\,r_p\,[\text{cm}])^2 \quad [\text{V}]. \tag{3.5}$$

Thus, high-density single-component plasmas require substantial voltages to confine them, unless r_p is small. In practice, when the space-charge potential of the plasma is not too high, the voltage applied to the end electrodes is typically $\sim 5 \times \Delta\phi$.

3.1.2 Radial confinement

The external electrostatic potential, ϕ_{ext}, applied for the longitudinal confinement of the plasma satisfies Laplace's equation $\nabla^2 \phi_{ext}(r, z) = 0$. Therefore, the axial potential well is a saddle point in the potential, since solutions to Laplace's equation cannot exhibit local minima. Thus, ϕ_{ext} cannot confine the particles radially. In addition, the space-charge of the non-neutral plasma is repulsive and pushes individual particles radially outwards. In the absence of external radial forces, the plasma will expand radially and particles will be lost on the walls of the cylindrical electrodes.

To counterbalance the radial expansion of the plasma, an axial uniform magnetic field, $\mathbf{B} = B\hat{\mathbf{z}}$, is applied in addition to the axial potential well. The rotation of the plasma around the axis of symmetry leads to a Lorentz force, $q\mathbf{v} \times \mathbf{B}$, pointing radially inwards. In simple terms, the Lorentz force balances the outward space-charge repulsion and the outward centrifugal force due to the rotation of the charge cloud, thus stopping particles from reaching the walls of the trap.

Cold-fluid theory of confinement

A first step in understanding the radial confinement of the plasma is to assume that the pressure forces due to the finite temperature of the plasma are negligible. Then, a necessary condition for confinement is given by the momentum balance equation [187]

$$mn (\mathbf{v} \cdot \nabla) \mathbf{v} = qn (\mathbf{E} + \mathbf{v} \times \mathbf{B}). \tag{3.6}$$

The radial component of eq. (3.6) is a statement of the balance between the outward centrifugal and electric forces and the inward magnetic force according to

$$q (E_r + v_\theta B) + \frac{mv_\theta^2}{r} = 0, \tag{3.7}$$

where v_θ is the equilibrium azimuthal velocity of a plasma fluid element. Similarly to the model in the previous section, the plasma is assumed to be infinitely-long with the rectangular radial density profile that led to eq. (3.2). Introducing the angular rotation velocity $\omega_r(r) = v_\theta(r)/r$, the solutions to the eq. (3.7) give the two cold-fluid rotation velocities

$$\omega_r = \omega_r^\pm = \frac{\omega_c}{2} \left[1 \pm \left(1 - \frac{2\omega_p^2}{\omega_c^2} \right) \right]. \tag{3.8}$$

Thus, the rotation velocity of the plasma column is obtained as a function of the cyclotron frequency, $\omega_c = qB/m$, and the plasma frequency, $\omega_p^2 = (nq^2)/(\epsilon_0 m)$. The solutions $\omega_r^\pm$ are independent of the radius r which shows that the plasma rotates as a rigid column about the axis of symmetry. ω_r^- (ω_r^+) corresponds to a slow (fast) rotation of the plasma column. For a low-density plasma or a high enough magnetic field, $2\omega_p^2/\omega_c^2 \ll 1$, the fast rotation frequency approaches the cyclotron frequency, $\omega_r^+ \approx \omega_c$, while the slow frequency ω_r^- corresponds to
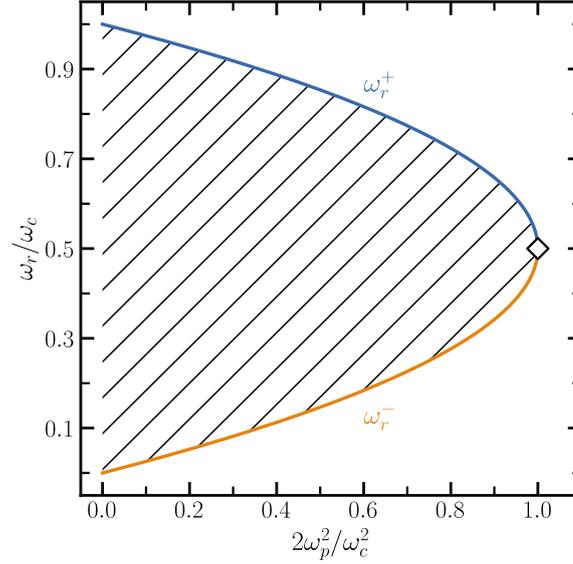
Figure 3.3: The cold-fluid angular rotation velocities ω_r^- and ω_r^+ as a function of the parameter $2\omega_p^2/\omega_c^2 \sim n/B^2$ (eq. (3.8)). The region above the fast rotation curve and below the slow rotation curve corresponds to the space-charge force being too large for radial confinement to be possible. The diamond marks the maximum plasma density that can be confined (Brillouin density n_B). The hatched area corresponds to the region of allowed rigid-rotor angular velocities for a plasma column in a finite-temperature thermal equilibrium (eq (3.19)) [190]

a slow $(E_r\hat{\mathbf{r}} \times B\hat{\mathbf{z}})$ rotation:

$$\omega_r^- \approx (qn)/(8\pi^2\epsilon_0 B). \tag{3.9}$$

Figure 3.3 shows the cold-fluid angular rotation velocities of eq. (3.8). It is important to note that the allowed values of the angular rotation velocity depend strongly on the parameter $2\omega_p^2/\omega_c^2 \sim n/B^2$, i.e. on the balance between the strength of the confining magnetic field and the plasma density. Whether the plasma is rotating with either the slow (ω_r^-) or the fast (ω_r^+) rotation velocity depends on how the non-neutral plasma forms inside or is injected into the trap. The two rotational equilibria described by fig. 3.3 were measured experimentally by [191] for a cold, uniform-density plasma column.

A second important characteristic of the cold-fluid rotation of the plasma can be extracted from fig. 3.3. For a fixed magnetic field, there is a maximum plasma density beyond which confinement is not possible. This maximum plasma density, n_B, corresponds to the condition $2\omega_p^2/\omega_c^2 = 1$ which is known as the Brillouin limit [192]. For an electron or positron plasma, the Brillouin density limit is

$$n_B = \frac{\epsilon_0 B^2}{2m_e} = 4.8 \times 10^{18}\,(B\,[\mathrm{T}])^2 \quad [\mathrm{m}^{-3}]. \tag{3.10}$$

However, special manipulation techniques must be used in order to compress the plasma all the way to the Brillouin limit. While this has been achieved in a laser-cooled ion plasma, the record so far for electron or positron plasmas is around 17% [187]. One interpretation of the maximum achievable density is given in [187] by rewriting eq. (3.10) as the ratio of the relativistic rest energy density of the plasma $n_b mc^2$ and the magnetic

energy density $B^2/2\mu_0$. Thus, the total stored rest energy of a single-component plasma is, at best, a fraction of the stored energy in the confining solenoid.

Finite temperature effects and thermal equilibria

A further step to examine the equilibrium properties of an axisymmetric non-neutral plasma column confined radially by a uniform magnetic field $B\hat{\mathbf{z}}$ is to remove the assumption that the plasma is cold. The description given below of the equilibrium state of the confined plasma is adapted from [187] and [190]. For a plasma at a finite temperature T, the equilibrium condition of eq. (3.6) is replaced by

$$mn \left(\mathbf{v} \cdot \nabla\right) \mathbf{v} = qn \left(\mathbf{E} + \mathbf{v} \times \mathbf{B}\right) - \nabla p, \tag{3.11}$$

where p is the plasma thermal pressure. As before, we assume that each fluid element is rotating with $\mathbf{v} = -\omega_r r\hat{\theta}$. Assuming that the plasma is ideal, the pressure is given by the ideal gas law, $p = nk_B T$. In general, an assumption about the heat flux is required to obtain a solution to the momentum equation. The assumption about the heat flux leads to an equation of state $pn^{-\gamma} = \mathrm{const.}$, where the value of γ depends on the exact assumption and on the isotropy of the energy distribution. Here, it is assumed that the heat transport is so rapid that isothermal condition ($\gamma = 1, T = \mathrm{const.}$) is valid. A necessary condition for the thermal equilibrium is that the temperature, T, is uniform within the plasma. Then, the axial component of eq. (3.11) can be written as

$$0 = -qn\frac{\partial \phi}{\partial z} - k_B T \frac{\partial n}{\partial z}. \tag{3.12}$$

Integrating this equation results in a Boltzmann-like relation between the density distribution and the electrostatic potential

$$n(r,z) = f(r)e^{-q\phi(r,z)/k_B T}, \tag{3.13}$$

where $f(r)$ is any function of radius. Thus, the plasma density collects in z at local minima of the electrostatic potential energy $q\phi$ [187]. The function $f(r)$ is determined by the radial component of eq. (3.11):

$$-mn\omega_r^2 r = qn\left(-\frac{\partial \phi}{\partial r} - \omega_r rB\right) - k_B T \frac{\partial n}{\partial r}. \tag{3.14}$$

Substituting n from eq. (3.13) and rearranging terms yields

$$\frac{k_B T}{f}\frac{\partial f}{\partial r} = -m\omega_r(\omega_c - \omega_r)r. \tag{3.15}$$

Thus, the radial density profile influenced by $f(r)$ is determined by the rotation frequency ω_r, or vice versa.

A special choice of f corresponds to the 'rigid-rotor' class of thermal equilibria for which the rotation frequency ω_r does not depend on r. An isolated non-neutral plasma column will relax to such an equilibrium state through binary collision processes [190]. In other words, the viscosity of the plasma acts to remove any

shears in ω_r. The relaxation of a non-neutral electron plasma column to global thermal equilibrium characterised by a rigid rotation has been experimentally verified [193].

Integrating eq. (3.15) and combining the resulting $f(r)$ with eq. (3.13) results in the density profile of a rigid-rotor plasma column in global thermal equilibrium [189]:

$$n(r, z) = C e^{-q\phi_{\text{eff}}(r,z)/k_B T},$$

(3.16)

where C is a constant of integration and the effective potential ϕ_{eff} is

$$q\phi_{\text{eff}} = \frac{1}{2} m \omega_r (\omega_c - \omega_r) r^2 + q\phi(r, z).$$

(3.17)

If we take $\phi(r = 0) = 0$ without loss of generality, C can be identified with the electron density on axis $(r = 0)$, $\hat{n}$. The first term in the effective potential is due to the rigid rotation through the uniform magnetic field which produces a radially confining potential partially offset by the centrifugal repulsion proportional to ω_r^2. The second term in the effective potential describes the longitudinal confinement ensured by the potential well along z created by the voltages applied to the end electrodes.

Equations (3.16) and (3.17) define the plasma density in terms of the total electrostatic potential. In turn, the potential is determined from the plasma density by Poisson's equation

$$\frac{\partial^2 \phi}{\partial r^2} + \frac{1}{r}\frac{\partial \phi}{\partial r} + \frac{\partial^2 \phi}{\partial z^2} = -\frac{n(r, z)}{\epsilon_0}.$$

(3.18)

Finally, eqs. (3.16) and (3.18) constitute a Poisson-Boltzmann system that can be solved uniquely [194] given the following plasma parameters: ω_r, T, C, and the boundary conditions defined by the geometry and the voltages of the confinement electrodes. This system of equations is highly non-linear. However, for the choice of a rigid-rotor equilibrium, it can be shown that the plasma column is radially confined, $n(r \to \infty) = 0$, provided that [190]

$$\omega_c \omega_r - \omega_r^2 - \omega_p^2/2 > 0.$$

(3.19)

Equation (3.19) is equivalent to the condition $\omega_r^- < \omega_r < \omega_r^+$, where ω_r^- and ω_r^+ are the cold-fluid rotation velocities defined in eq. (3.8). The region of allowed rigid-rotor angular velocities for a plasma column in thermal equilibrium is also shown in fig. 3.3. Equation (3.19) is also a necessary and sufficient condition for the density profile $n(r)$ to be a monotonically decreasing function of r [190]. Therefore, a confined plasma with an increasing radial density profile tends to be unstable.

Section 7.1.3 presents an iterative numerical method to solve the Poisson-Boltzmann system fully in order to compute the density distribution of a cylindrical confined plasma in global thermal equilibrium. However, some important qualitative features can be highlighted here. If the rotation velocity ω_r is significantly different to the cold-fluid angular velocities ω_r^- and ω_r^+ (see fig. 3.3), the radial density profile has a bell shape and extends
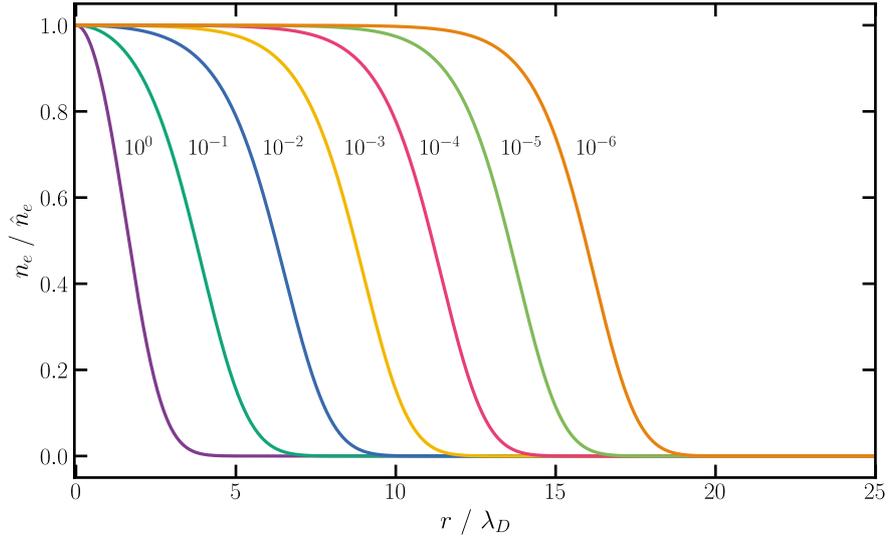
Figure 3.4: The radial density profile $n_e(r)$ normalised to the maximum value $\hat{n}_e$ for a confined non-neutral electron plasma in global thermal equilibrium. The curves represent the self-consistent solutions to the eqs. (3.16) and (3.18) for an infinitely long plasma column. Each curve corresponds to a different rigid-rotor angular velocity $\omega_r = \omega_r^-(1 + \epsilon)$, where the values of ϵ are indicated in the figure. The other relevant plasma parameters are: $T = 1\,\mathrm{eV}$, $\hat{n}_e = 1 \times 10^{15}\,\mathrm{m}^{-3}$, $B = 30\,\mathrm{mT}$, $\lambda_D = 2.4\,\mathrm{mm}$.

over a few Debye lengths. The Debye length is defined as $\lambda_D = \sqrt{(\epsilon_0 k_B T)/(\hat{n} e^2)}$, where $\hat{n}$ is the maximum electron density at the centre of the plasma column. On the other hand, if ω_r is closely tuned to either ω_r^- or ω_r^+, the density profile $n(r)$ is approximately uniform over many Debye lengths out to the radial edge of the plasma column where $n(r)$ drops abruptly to zero over a distance of a few times λ_D [190]. This characteristic radial density distribution is shown in fig. 3.4. Furthermore, the confined rotating thermal equilibrium has a Maxwellian velocity distribution $f_{\mathrm{eq}}(r, z, \mathbf{v})$.

An alternative explanation of the mechanism that leads to the radial confinement of the plasma is based on the conservation of the constants of the motion. For example, the total plasma angular momentum P_θ is a constant of the motion given by [187]

$$P_\theta = \sum_i^N p_{\theta i}(t) = \int d^3 r d^3 v f_{\mathrm{eq}}(\mathbf{r}, \mathbf{v}, t)(m v_\theta r + q B r^2 / 2), \tag{3.20}$$

where $p_{\theta i}$ is the azimuthal component of the single-particle canonical angular momentum. For a sufficiently large magentic field, the mechanical part of $p_{\theta i}$ is negligible compared to the contribution of the vector potential, resulting in

$$P_\theta \cong \frac{qB}{2} \int d^3 r d^3 v f_{\mathrm{eq}} f(\mathbf{r}, \mathbf{v}, t) r^2 = \frac{qNB}{2} \langle r^2 \rangle, \tag{3.21}$$

where $\langle r^2 \rangle$ is the mean-square plasma radius and N the total number of particles. Since P_θ is a constant of the motion, the mean-square plasma radius remains constant during the evolution of the plasma. The practical consequence of this result is that good plasma confinement over long periods of time requires the minimisa-

tion of the effect of any non-conservative forces, such as the collisional drag on the neutral background gas. Furthermore, the alignment of the external magnetic field with the axis of symmetry of the trap electrodes is equally critical [195] to avoid an external force with a component along $\hat{\theta}$. For example, it has been shown [196] that a tilt in the magnetic field that creates a transverse field component that is a few parts in a thousand of the longitudinal component can lead to variations in the maximum confined plasma density larger than 100% [196].

In practice, confined plasmas expand slowly due to imperfections in the trap or due to collisions with the neutral gas in a phenomenon which is termed 'outward plasma transport'. Several scaling laws for the expansion rate have been determined experimentally [188, 197, 198]. In general, the constants involved in these laws are dependent on each individual device that traps the plasma under study. When radial transport is present, angular momentum and energy are delivered to and lost from the plasma. Whether a steady-state plasma will be near thermal equilibrium at low enough temperature for the Debye length to be much smaller than the radius of the column remains a matter of experimental investigation in each plasma-confining apparatus.

3.1.3 Charged particle beam focusing

Due to the radial electric field produced by a long, cylindrical cloud of electrons, a Penning-Malmberg (PM) plasma trap was proposed as a focusing device for positively-charged particle beams [49]. The linear dynamics of a single positively-charged particle is derived in appendix B for the case of an infinitely-long electron-plasma column with a uniform radial density distribution. This simplified assumption on the distribution of the electron density is not far from the more realistic case of a fully thermalised confined plasma.

In the longitudinal direction (z-axis), the electrons rearrange themselves to completely screen the external electrostatic potential generated by the cylindrical electrodes of the trap. Thus, the electron density drops down along the axial direction from the maximum value in the bulk of the plasma to zero at the two ends of the plasma. The variation of the electron density at the ends of the plasma takes place over a region on the scale of the Debye length, λ_D. In a typical PM trap, the length of the plasma is much larger than the Debye length ($L \gg \lambda_D$), and, thus, the focusing effect inside the uniform region of the plasma is considerably more significant than that experienced by a beam when entering or exiting the plasma. Furthermore, the edges of the plasma are not straight lines as the electron distribution is dictated by the curved field lines produced by the cylindrical electrodes. For well-confined electron plasmas [194], the length of the electron cloud is maximum on the central axis and decreases radially as is noticeable in fig. 3.6. However, the variation of the length of the plasma at different radial locations also arises on the scale of a few Debye lengths.

In the transverse plane, it was shown that a globally-thermalised plasma exhibits a central region of uniform density and a sharp radial drop off over a distance of a few Debye lengths (see fig. 3.4). Thus, if the radius of the plasma is larger than the radius of the beam envelope, all the particles in the beam pass through the uniform density region of the electron cloud. An example of the radial electric field produced by a confined electron plasma in the mid-plane of the trap is shown in fig. 3.5. Inside the region of uniform electron density, the strength of the electric field increases linearly with the radius. As a result, all the positively-charged particles passing
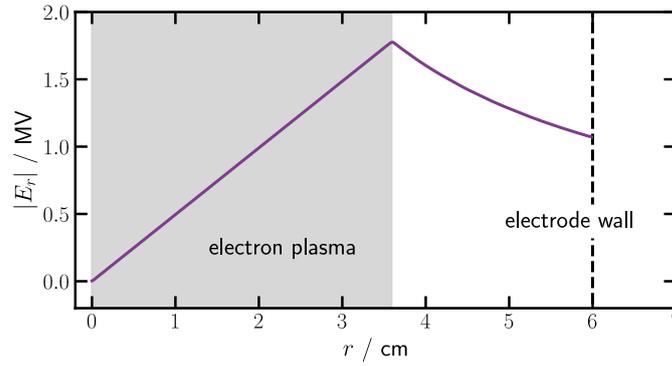
Figure 3.5: The magnitude of the radial electric field in the midplane of a non-neutral uniform electron plasma confined inside a cylindrical anode. The grey colour indicates the region occupied by the plasma. The field terminates at a radius of $6\,\mathrm{cm}$ on the conducting surface of an electrode.

through the plasma experience an ideal focusing force in both transverse planes simultaneously. The particles with the same kinetic energy are brought to a single focus downstream of the plasma as shown in fig. 3.6. However, outside of the plasma, the radial electric field starts to decrease non-linearly. Hence, the particles that pass through the lens outside of the plasma are subject to non-ideal focusing. Figure 3.6 shows that the protons passing partially or fully outside of the plasma can be seen to end up further downstream of the common focal point of the protons that pass completely through the electron cloud.

Furthermore, fig. 3.6 shows the chromatic dependence of the focusing strength of an electron plasma on the kinetic energy of non-relativistic protons. An overview of the beam-focusing capabilities of an electron plasma lens (eqs. (B.17) and (B.24)) with respect to the fields required to confine the electrons in the longitudinal and radial direction within a PM trap (eqs. (3.5) and (3.10)) is shown in fig. 3.7.

'Operation function' of a plasma lens

In the model of the plasma confinement presented above, the strength of the external fields of the plasma trap determine the maximum electron density than can be successfully contained within the volume of the trap. However, two upper limits for the plasma density are found as a result of applying the condition for radial confinement separately to the condition for longitudinal confinement. The strength of the axial magnetic field imposes an upper limit on the plasma density according to eq. (3.10) to prevent electrons from being lost radially. Simultaneously, the value of the voltage difference applied between the end gates and the rest of the electrodes in the trap is required to be higher than the space-charge potential of the plasma described by eq. (3.4) to stop electrons from being lost longitudinally. In practice, it was found that the magnetic field, B_z, in the trap and the voltage difference applied to the electrodes, V_A, need to be closely matched in relation to one another to produce highly homogeneous plasma columns and to avoid the formation of plasma instabilities. Hence, a 'work function' or 'operation function' can be defined for a plasma lens in the form $V_A = V_A(B_z)$ for the field settings which allow a stable plasma to be confined. The ideal configuration of the fields defined by the operation function is the one that leads to equal limits on the maximum plasma density from the conditions of radial and longitudinal
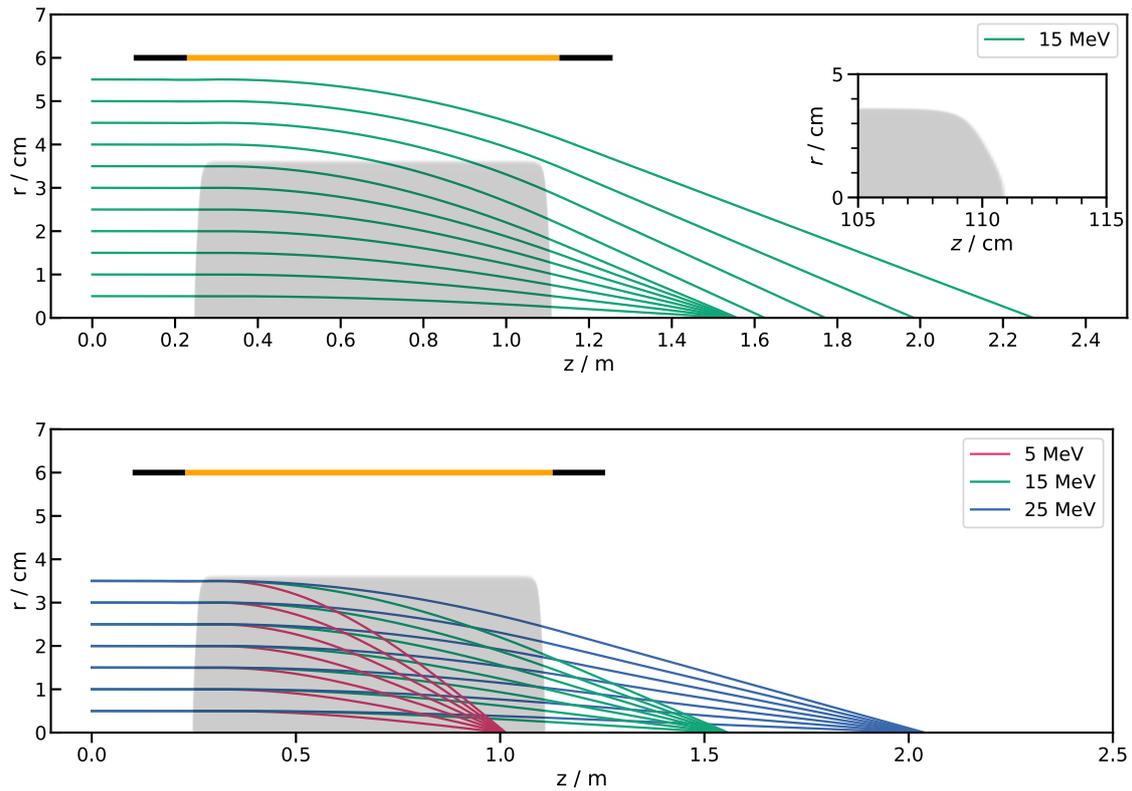
Figure 3.6: Ray tracing through a non-neutral electron plasma confined inside a Penning-Malmberg trap. Each ray represents the trajectory of a proton starting at a different initial radial position with zero transverse momentum. The grey area indicates the electron density ($n_{\mathrm{e,max}} = 5.4 \times 10^{15}\,\mathrm{m}^{-3}$) inside the trap obtained as the self-consistent numerical solution of eqs. (3.16) and (3.18) which was used as a field map for particle tracking with GPT [199]. The position and size of the central electrode (orange) and the end electrodes (black) is also indicated. **Top:** Only the particles passing through the plasma experience an ideal linear focusing force that brings all the particles to the same focal point (see also fig. 3.5). The inset shows the scale of the curvature of the plasma at one of its ends. **Bottom:** The focusing strength of a non-neutral electron plasma depends on the kinetic energy of the beam passing through the lens.
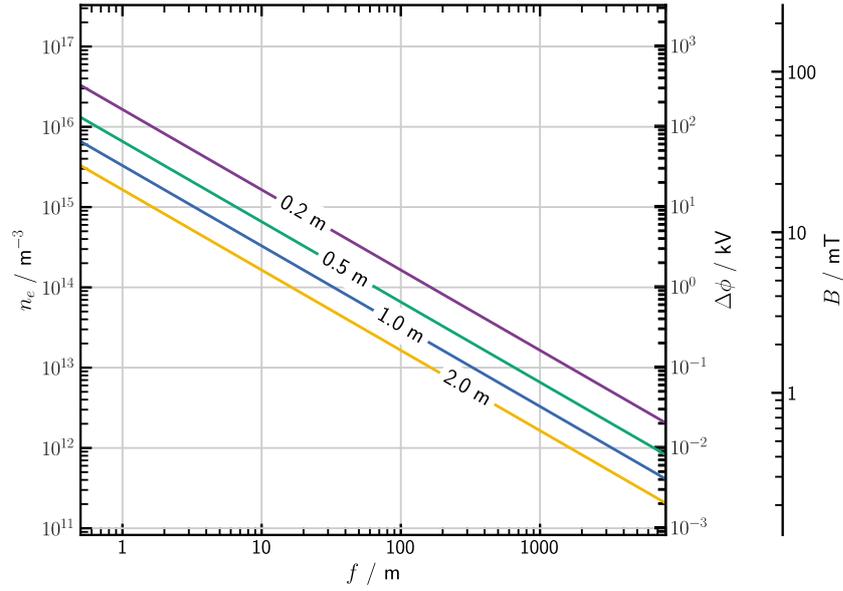
Figure 3.7: The density n_e of a non-neutral electron plasma as a function of the thin-lens focal length that is produced by the space-charge cloud for a beam of $15\,\mathrm{MeV}$ protons. The corresponding length of the plasma is indicated on each line and the radius is $r_p = 3\,\mathrm{cm}$. The secondary axes on the right show the minimum strength of the fields that can confine a plasma with a specific density on the vertical axis on the left. A minimum voltage $\Delta\phi$ must be applied to the end electrodes of the trap and a minimum axial magnetic field B has to be provided by the trap solenoid.

confinement [46]. Combining eqs. (3.4) and (3.10), the plasma lens work function is

$$V_A = \frac{q r_p^2}{8 m_e} \left[1 + 2\ln\left(\frac{r_w}{r_p}\right) \right] B_z^2. \tag{3.22}$$

Even though this curve in the parameter space (V_A, B_z) corresponds to idealised plasma-confinement conditions, experimental tests of space-charge lenses [51] showed that the most efficient filling of the lens with plasma is achieved when the confining fields are configured according to eq. (3.22) along a similar operational curve as the one shown in fig. 3.8.
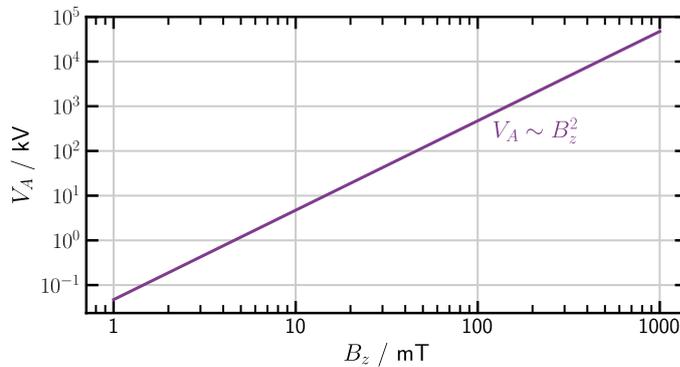
Figure 3.8: The operation function of eq. (3.22) for the idealised configuration of the voltage V_A and the magnetic field B_z of an electron plasma trap in the absence of an ion beam. The plasma has a radius $r_p = 3\,\mathrm{cm}$ and the radius of the trap wall satisfies $r_w = 4 r_p$.

Previous numerical simulations and electron density measurements suggested that plasma instabilities appeared in several versions of the Gabor lens when the parameters of the lens are tuned away from its work function [200, 201]. One of the most common instabilities in a low density non-neutral electron plasma is the diocotron instability [202, 203]. It was first observed in low-current annular electron beams [204], but it is also typical for electron plasma columns confined radially by a uniform axial magnetic field B. Its driving mechanism is an azimuthal shear in the angular $\boldsymbol{E} \times \boldsymbol{B}$ rotation velocity, $\omega_E(r) = -E_r(r)/rB$, of the plasma. Here, $E_r(r)$ is the equilibrium radial electric field. The instability is important in strong magnetic fields for which $\omega_{pe}^2/\omega_{ce}^2 \ll 1$. While at the Brillouin flow limit, which gives the radial confinement condition, the ratio $\omega_{pe}^2/\omega_{ce}^2$ is $1/2$, stronger magnetic fields may be required in an experimental setup, thus lowering this ratio well below 0.5.

Numerical simulations show that the diocotron instability is linked to annular electron profiles. If the plasma is regarded as a 'tube' of electrons, the instability can be explained as the interaction between two waves that propagate azimuthally in opposite directions on the inner and outer surfaces of the plasma 'tube'. An additional factor believed to promote the instability is the residual gas pressure as electrons loose energy in the collisions with residual gas molecules [205]. For the interested reader, a further discussion of the conditions that can lead to the onset and growth of the diocotron instability is included in appendix D. In practice, a plasma lens has to be filled as homogeneously as possible in order to operate in a regime away from the diocotron instability.

3.2 Development of non-neutral electron plasma lenses

3.2.1 Historical overview

In 1947, the use of a cylindrical electron plasma as a focusing space-charge lens for high-energy ion beams was suggested by D. Gabor [49]. In the original proposal, the electron plasma is confined in a setup very similar to the Penning-Malmberg trap by a central anode at high positive potential with respect to two end 'guard electrodes', and by two magnetic coils of unequal lengths located around the electrodes. The two coils are wound in opposite direction to generate a cusp in the magnetic field. A hot cathode in the form of a circular loop is placed near the cusp to fill the lens with electrons until the axial external electrostatic potential is depressed to nearly the cathode potential and equilibrium is reached. The magnetic cusp was part of the design to ensure that the electrons are injected into the trap with a component of momentum transverse to the magnetic field lines. In the absence of any transverse momentum, the electrons spiral closely around the field lines through cyclotron motion and are prevented from crossing magnetic flux lines and filling the region of the trap near the symmetry axis. In the initial paper [49], Gabor also mentioned the theory of Brillouin [192] that specifies the maximum electron density that can be confined with the proposed trap, eq. (3.10), and highlighted the advantage of an electron plasma lens over both magnetic and electrostatic lenses for high energy ion beams.

Several years after the initial proposal, numerous attempts were made at the major laboratories around the world to create a practical electron plasma lens based on the original design of Gabor. The initial experimental

reports confirmed that the stable operation of a space-charge lens is possible and several configurations of the coils and electrodes were tested [40]. However, the longer lenses had to be operated in tens-of-milliseconds-long periods to avoid the formation of continuous high-density glow discharges and large time variations of the focusing strengths were common. Even from the initial experiments, it was clear that a low background pressure ($<10^{-5}$ mbar) and specific electric and magnetic field configurations are required to obtain a stable dynamic equilibrium after the initial formation of the plasma. Given these operational challenges, ion beams were successfully transported [40, 41] through the lenses and the measured focal lengths were a significant fraction of the theoretically predicted ones [40]. Further tests of the theoretical model of the space charge confirmed the approximate linear dependence of the focal length on the voltage applied to the trap electrodes and on the kinetic energy of the ion beam [41]. Furthermore, in specific cases, the focusing force was shown to be linear across the lens aperture to within a few percent [40]. A measurement of the transported current and angular acceptance proved the advantage of a space-charge lens in comparison to a magnetic quadrupole lens for a simple beam transport system [41].

One particularly great achievement in terms of electron plasma confinement for beam focusing is reported in [42] where the use of a 65 cm-focal-length plasma lens is described for the successful focusing of multiple ion species with energy of a few MeV. The aim of the lens was to form an ion microbeam for the probing of targets by scanning the microbeam over their surface. Ion-beam spot sizes as low as 5 µm were obtained [206]. In order to achieve such small spot sizes, the metre-scale focal length was obtained with a plasma trap consisting of nine co-axially mounted electrode rings maintained at voltages up to several kV that were able to confine an electron column with a length of about 16 cm. Moreover, the blue light emitted axially by the plasma was analysed with a spectrograph to extract the radial distribution of plasma [207]. Separately, the deflection of a proton beam with energy $\sim$MeV was measured at various radial locations to determine the absolute electron density. The plasma column under investigation had a nearly constant density in the range 10^{15}–10^{16} m^{-3} inside a radius of several millimetres [207].

As a result of the previous reports of excellent optical quality and the strong focusing strength achievable with a space-charge lens, a set of measurements at Fermilab [43] tried to determined whether a plasma lens can be used to match a 30 keV, 30 mA proton beam into a radiofrequency quadrupole (RFQ). It was well-known that intense low energy ion beams suffer from space-charge-induced emittance growth and the aim was to verify whether a plasma lens can provide neutralisation of the beam space charge while simultaneously focusing the beam to the small spot size necessary for matching into an RFQ. Stable operation of the lens has been observed for periods of 24 hours for a maximum applied voltage of 12 kV and the beam transmission was measured to be 100% [43]. However, measurements of the beam emittance upstream and downstream of the lens showed emittance growth by a factor of three to four. Part of this emittance growth was attributed to the self-repulsion of the particles within the beam in the drift spaces outside of the lens, but it was not clear whether lens aberrations also contribute to the emittance growth.

It is important to note that the experimental efforts described above from before 1990 involved devices

based upon the ionisation of background-gas atoms present within the apparatus to form discharge plasmas. In practice, the discharge was simply started by turning on the power supplies for the magnet and the high voltage on the anode or the end electrodes [208]. After some initial out-gassing the discharge was observed to become quasi-stable. To achieve a stable electron cloud, it was found useful to lower the anode potential to below the value necessary to initiate the cold-cathode Penning discharge [209]. It was often the case that the transition of the plasma from a chaotic state to a steady state was dictated by 'chance' [206]. On the one hand, the steady-state plasmas that some devices supported were self-stabilising. On the other hand, the stabilisation mechanism was observed to accommodate more than one plasma profile for identical confining fields [206]. Soon, it was realised that an electron source may be necessary in order to produce a plasma column that has a radius comparable to the ion beam size and, thus, can provide linear focusing. Simulation studies at Fermilab suggested that two space-charge lenses were strong enough to match an ion beam into an RFQ even though the lenses were limited to about $10\,\text{kV}$ electrode voltage due to sparking [208]. Nonetheless, the beam reached a radius of nearly $3\,\text{cm}$ in the simulations which had to be matched by the radius of the plasma to achieve ideal focusing.

Also at the beginning of the 1990s, a programme of investigations of electron space-charge lenses started in Frankfurt [44]. The initial measurements were done on a $5\,\text{kV}$ lens that relied on the same electron production mechanisms as its earlier predecessors: gas discharge. It was further highlighted in the initial report [44] that there seemed to be no method to control the production density in different regions inside the lens. Numerical simulations suggested that electrons move inside the lens from regions of density higher than the radial limit for trapping into regions with lower density, thus resulting in a final homogeneous plasma distribution. Beam transport experiments showed that the lens could produce a focal length of $20\,\text{cm}$ by focusing of a $10\,\text{keV}$ He beam with negligible emittance growth. However, the focal length measured experimentally was a factor of four larger than expected, indicating that the lens was only filled only to about 25% of the radial trapping density limit. Furthermore, a measurement of the energy of the ions escaping the lens on the axis provided information about the electric potential distribution in the lens from which a lower limit of 40% was calculated for the degree of lens filling. Amongst the conclusions of the operation of this initial space-charge lens in Frankfurt, focusing aberrations, emittance growth, and non-homogeneous lens filling were found to be the main problems in order to advance the lens from being an experimental apparatus to a reliable ion-beam-focusing device. The development programme started in Frankfurt to produce an operational space-charge lens has continued [210]. A short review of some of the experimental results from the research programme in Frankfurt is given in sub-section 3.2.2. The interested reader can find in appendix A an overview of additional attempts to create a practical plasma lens for focusing heavy ions and negative ion beams.

3.2.2 Beam transport and plasma stability

The experimental efforts after 1990 to develop a working space-charge lens have been mainly focused on improving the focusing quality of the plasma and understanding and reducing any source of beam emittance

growth. Any non-linearities observed in the focusing of ion beams were typically attributed to plasma instabilities or a non-homogeneous electron distribution inside the lens. It is important to highlight again that most of these efforts involved devices that produced weakly-confined quasi-steady-state electron plasma by the uncontrolled ionisation of the background gas. Some of the most important experimental results of ion transport through weakly-confined plasmas are summarised in this subsection.

Degree of lens filling

As mentioned previously, probably the most substantial R&D programme on space-charge lenses has been carried out at Goethe University Frankfurt. Several prototypes of electron-plasma lenses were designed, built and tested [48] with lengths from several centimetres up to $2\,\mathrm{m}$, radii in the range 5–$10\,\mathrm{cm}$, and maximum electrode voltages up to $65\,\mathrm{kV}$ [53, 211]. For one of the lens designs [53], the focusing capability was proven by the use of a $1.2\,\mathrm{mA}$, $440\,\mathrm{keV}$ $\mathrm{He^+}$ beam for a plasma with an estimated average density of about $10^{14}\,\mathrm{m^{-3}}$. In a first step, the average electron density was deduced by measuring the focal length of the lens. Secondly, the average electron density was compared to the predicted value obtained from 2-D hydrodynamics and 3-D particle-in-cell codes to estimate the plasma-trapping efficiency of the apparatus. The trapping efficiency reached $\sim 50\%$ at low anode voltages and gradually degraded for higher magnetic or electrostatic fields to about 20%. It was found that the maximum trapping efficiency is obtained for a specific anode voltage for each fixed value of the strength of the magnetic field. The line that gives the maximum trapping efficiency in the 2-D parameter space of anode voltage vs. magnetic field was called the 'operation function' of the lens. At the same time, the work function also represents the configuration of external fields that leads to a homogeneous electron distribution inside the lens.

The extent to which a space-charge lens can be filled with a plasma close to the maximum electron density predicted by theory was also the subject of a more recent set of experimental measurements and numerical simulations [46]. The 2-D distribution of electrons inside a particular configuration of the lens was obtained numerically by an iterative scheme that adjusts the local plasma density until the local radial electric field required for the local equilibrium of a fluid element becomes consistent with the electric field generated by the space-charge density of the plasma. Thus, the calculated local space-charge densities satisfy the radial confinement conditions. In order to couple the different transverse planes along the longitudinal direction of the trap, the numerical scheme assumes a full thermalisation of the plasma along the magnetic field lines according to Botzmann's distribution. As a result, the algorithm produces the equilibrium space-charge density in the lens for given geometry of the external magnetic and electric fields and for an electron temperature provided as an input. For the geometries and confining fields simulated in [46] for two different lenses, the average electron densities obtained from the numerical scheme were reported as $\sim 50\%$ and $\sim 60\%$ of the maximum electron densities predicted by the theory of plasma confinement in the radial and longitudinal directions, respectively. Looking at the particular density distributions inside the two lenses modelled in [46], nearly linear focusing was predicted from a homogeneous region near the axis of the plasma with density variations of up to ± 2–3%

for beam radii smaller than $15\,\mathrm{mm}$. However, the beam emittances measured downstream of the lens were significantly higher than those predicted by simulations, although comparable to the emittances produced by conventional lens systems such as solenoids or Einzel lenses.

In parallel to the numerical simulations, the work of [46] also studied the transport of He beam through the two space-charge lenses tested in the experimental setup. Measuring the focal length of each lens allowed the determination of the average plasma densities produced in the lenses during the tests. The data showed lens-filling factors around $40\%/50\%$ for one lens and $50\%/60\%$ for the second relative to the maximum densities predicted by the theory of radial/longitudinal confinement, respectively. These values matched very well with the predictions from the numerical simulations. It must be noted that these relatively high filling factors that were reported in [46] correspond to absolute electron-plasma densities in the range $0.8\text{–}3 \times 10^{14}\,\mathrm{m}^{-3}$.

Onset of plasma instabilities

It was found that the confining fields need to be set close to the operation function in order for the ion beam passing through the lens to exhibit negligible emittance growth. A large decline in the beam quality was attributed to plasma instabilities. One particular example that was reported is that of a hollow electron column which then leads to plasma waves propagating azimuthally within the electron ring [53]. The radial profile of such an unstable plasma was detected by optical imaging triggered by an oscillation of the measured current of ions that leave the lens on axis. The fact that non-uniform electron clouds cause significant aberrations was reported even from the early investigations of ion-beam transport through a space-charge lens. For example, focusing aberrations were described [52] as a result of the diocotron instability which appears when the electron density near the surface of the electron cloud is higher than in the central zone. Such a density profile is created when the electrostatic potential of the plasma closely approaches the external electrostatic potential generated by the trap electrodes. Hence, a lens should be operated at higher voltages and moderate magnetic fields to decrease the aberrations from the diocotron instability. The diocotron instability was also identified as a possible cause of the hollow plasma columns imaged in [53]. Numerical simulations confirmed the hypothesis that the diocotron azimuthal instability rises when the lens is filled with electrons over the maximum confinement density limit [201]. Further measurements showed a dependency of the plasma state not only on the enclosing fields, but also on the residual gas pressure, which was observed to be another factor that can induce the diocotron instability [200]. For the same confining fields, an increase in the residual gas pressure led to the formation of hollow plasma columns. The mechanism behind the dependency of the plasma state on the residual gas pressure might involve the production rate of electrons from ionisation events within the volume of the lens and the localisation of the more intense electron-production regions.

In 2012, the performance of a different lens prototype was tested [47, 51] at the GSI High Current Injector to evaluate a possible application of the lens for the transport of intense ion beams which, at low energies, propagate under a strong influence of the beam space-charge. One major issue highlighted in [51] was the generation of an homogeneous electron column with a radius comparable to the relatively large expected beam radii of up to

$50\,\mathrm{mm}$. Hence, nonlinearities in the focusing capabilities of the lens were observed due to plasma instabilities or the Debye density drop-off at the edge of the plasma cloud. In order to investigate the extent of the nonlinear effects, beam-transport measurements were conducted in two parts which are detailed below.

Emittance-dominated beam transport

In a first set of experiments, the capabilities of the space-charge lens described in [47, 51] were tested by transport of a $3\,\mathrm{mA}$, $50\,\mathrm{keV}$ He$^+$ beam. The low current of the beam implies that the influence of the beam space-charge on the transport is negligible, and, thus, an electron plasma acts as a focusing device. The emittance measurements downstream of the lens were matched well by beam-tracking simulations in terms of the orientation and shape of the transverse phase-space distribution of the beam [47]. However, the simulations overestimated the electron density in the lens and the magnetic field had to be tuned to improve the matching [51] to the measurements. Similarly, the emittances obtained from simulations were consistently much lower than the measured ones. Nonetheless, the simulations proved to be a very useful tool to identify the distribution of the plasma in the lens that lead to a specific shape of the beam in the transverse phase-space after passing through the lens. Hollow electron columns were associated with major aberrations and large emittance growth. At high magnetic fields close to the work function of the lens, the density became more homogeneous and the emittance was better kept under control. At magnetic fields above the operation point set by the lens voltage, the plasma was observed to be highly unstable. The conclusion of the experiments with the He beam was that the setup of the strength of the confining fields in relation to the operation function defined by the plasma confinement theory is of great importance to ensure a homogeneous electron density and the stability of the plasma. On the technical side, high-voltage breakdowns were observed as a result of metallic spots on the insulators and the presence of a metallic splint. Thus, both the presence of impurities during installation and material fatigue can impact the operation of the lens at high voltages.

Space-charge-dominated beam transport

A second category of experiments reported in [47, 51] investigated the transport of a $3.1\,\mathrm{keV/u\,Ar^+}$ beam with a high current of up to $35\,\mathrm{mA}$. With increasing ion current the beam transport starts to be affected by the space charge of the beam. At the same time, the ion beam acts as a source of secondary electrons by ionising the residual background gas or by beam losses on the vacuum pipe. Thus, the passing of the beam influences the degree to which the lens was filled by changing the plasma distribution from the one that is maintained by the confining fields in the absence of the ion beam. For the Ar$^+$ beam used during the measurements, it was estimated [51] that at residual gas pressures around $10^{-6}\,\mathrm{mbar}$ and beam currents above $20\,\mathrm{mA}$, one $1.25\,\mathrm{ms}$ ion-beam pulse can lead to the production of a number of electrons that is only a factor of two or three smaller than the total number of electrons trapped inside the lens in the absence of the beam. Qualitative evidence for the production of secondary electrons by the ion beam was also obtained in these measurements. The focusing strength of the lens was measured to be higher for larger beam currents when the confining fields and the residual
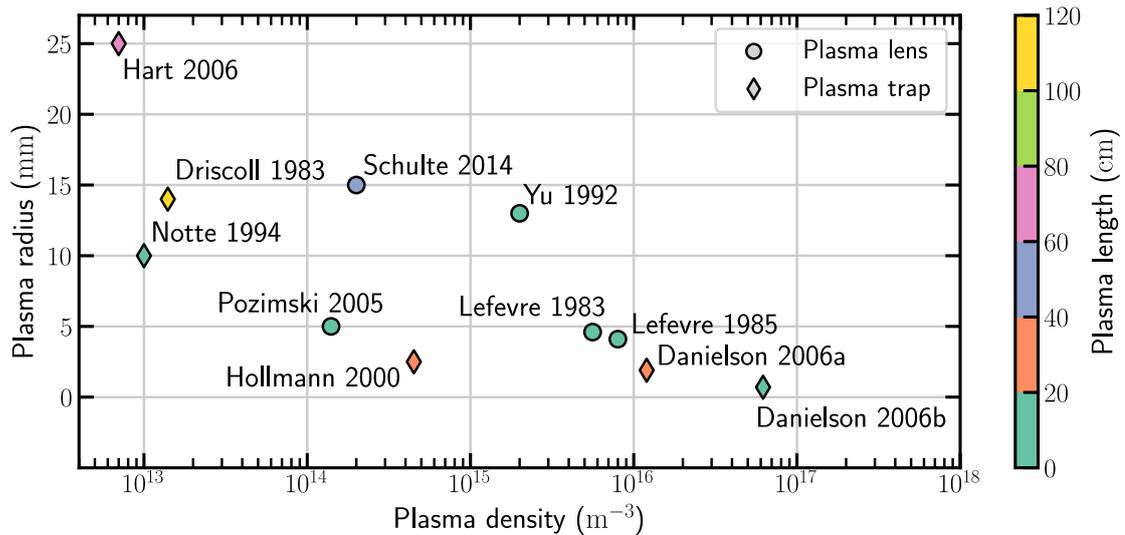
Figure 3.9: A selection of measurements of typical plasma density, radius and length that were reported for plasma-confining devices. The apparatuses that were specifically designed to act as space-charge lenses and that were showed to focus ion beams are termed 'plasma lenses'. The examples categorised as 'plasma traps' refer to devices intended entirely for electron or positron plasma confinement. Data taken from: Driscoll 1983 [197], Lefevre 1983 [42], Lefevre 1985 [206], Yu 1992 [52], Notte 1994 [195], Hollmann 2000 [212], Pozimski 2005 [46], Hart 2006 [213], Danielson 2006a [214], Danielson 2006b [215], Schulte 2014 [51].

gas pressure were kept constant. An alternative explanation of this observation is the reduction of the electron loss mechanism due to the positive electrostatic potential created by the ion beam near the axis of the lens.

When the lens parameters were varied for a fixed ion-beam current, the focusing strength was seen to increase with the magnetic field up to an upper threshold at which plasma instabilities were observed together with the associated beam emittance growth [51]. Nonetheless, the focusing effect of the electron plasma was successfully demonstrated with metre-scale focal lengths. For example, a parallel Ar^+ beam was achieved from a diverging one with negligible emittance growth. The confinement efficiency achieved in these experiments was approximately 14%.

3.2.3 Discussion

The overview presented in this section shows that there have been numerous attempts to create a practical space-charge lens based upon the original Gabor design. The majority of the devices were capable of producing a plasma and of focusing an ion beam. However, the homogeneity and stability of the plasma, and, thus, the quality of the focusing that was observed were typically below those required by a lens that can be operated reliably in an accelerator beam line. The plasma parameters of a selection of the lenses that were investigated more extensively in beam transport experiments are shown in fig. 3.9. It is important to note that the majority of space-charge lenses built and tested so far relied on the spontaneous formation of the plasma by the ionisation of the background-gas atoms. Under the appropriate configuration of the confining fields of the trap, the resulting plasmas were typically weakly-confined and their dynamic equilibrium and density distribution were difficult to

control. Nonetheless, sufficiently high electron densities were achieved to satisfy the requirement of relatively short focal lengths for high-energy ion beams.

In addition to the plasma-confinement devices that were built to act as a focusing lens for ion beams, fig. 3.9 shows the parameters of a selection of plasma traps that were designed and built for the specific purpose of studying the properties of confined electron or positron clouds. The low-energy, non-neutral plasma community has a long history of very successful efforts to investigate magnetically-confined non-neutral plasmas and develop the corresponding plasma traps [187, 216]. In order to trap electrons or positrons successfully at high density or for long periods of time, specific techniques have been developed to address the issues that negatively impact the stability of well-confined plasmas [187, 195, 197, 217–219]. A few important examples of the issues that have been separately addressed are: plasma densities at large fractions of the Brillouin limit, high background-gas pressures, high plasma-to-wall radii ratios, and long plasma lengths. All these factors contribute to the reliable operation of a space-charge lens. Thus, there is a strong rationale for the techniques developed for electron- and positron-plasma confinement to be integrated into the development of a practical non-neutral plasma lens.

In the first three chapters of this thesis, the principles of laser-plasma acceleration were described and the motivation for using laser-driven proton and ion sources for radiobiological research and radiotherapy as well as the challenges related to the capture and energy selection of the particles were explained. Moreover, the basic concepts of non-neutral plasma confinement and proton- and ion-beam focusing were introduced, and the most important previous attempts to build a working plasma lens were discussed. The following chapters will present the work that I performed during my PhD, starting with an investigation of the anomalous focusing effect observed with the first prototype of a Gabor lens built at Imperial College London.

Chapter 4

The Imperial College plasma (Gabor) lens

A prototype Gabor lens was built at Imperial College London to generate high-density electron plasmas with similar beam-focusing capabilities as those required for LhARA. The prototype was tested with beams of protons to investigate the focusing properties of the lens. The beam test resulted in limited experimental success due to a deleterious focusing effect which was recorded in several separate measurements. The non-ideal focusing effect was not predicted by the previous considerations of the lens design and of the mechanism by which the plasma was produced within the lens. Based on the observations from the beam test, simulations of the prototype lens and the associated electron plasma were performed to investigate possible causes of the anomalous focusing effect.

This chapter starts with a presentation of the configuration of the prototype lens (section 4.1). Then, the results of the simulations of the prototype lens producing a stable plasma (section 4.2) are presented. Additional simulations of several distinct plasma instabilities are described and their results are compared to the experimental observations from the beam test of the prototype lens (section 4.3). Lastly, the conclusions drawn from the simulation results are discussed (section 4.4). All PIC simulations presented in this chapter were performed with VSim [244] version 10.0.0.

4.1 First prototype of a Gabor lens built at Imperial College London

Along with the rapidly growing interest in the development of laser-accelerated proton and ion beams, the advantages of focusing systems based on a non-neutral plasma have been identified. A compact optical system comprising of Gabor space-charge lenses was proposed to collect, focus, and energy-filter laser-produced protons [186, 220, 221]. The proposed system is more compact and cost-effective than conventional alternatives for beam focusing and passive momentum selection. Simulations of beam-lines consisting of Gabor lenses and RF cavities showed that increased radiation doses can be delivered from a laser-driven source. In particular, a system based on space-charge lenses can be kept compact and reduce the overall cost of a facility more substantially for proton beams with energies suitable for the treatment of patients [66].

As a result of the validation of the proposed beam-lines in beam-tracking studies and of the foreseen
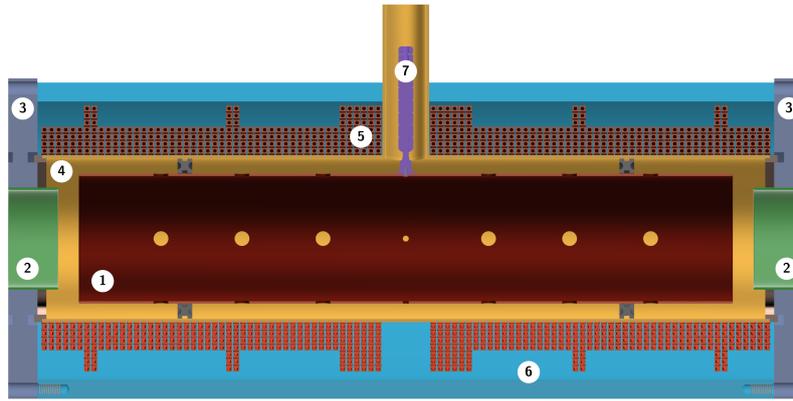
Figure 4.1: Internal structure of the IC Gabor lens viewed in longitudinal cross-section. The main components are: 1-central anode, 2-end electrodes, 3-end flanges, 4-vacuum tube, 5-pancake coils, 6-outer tube, 7-high-voltage feed-through. Diagram taken from [54].

advantages of the space-charge lenses, a first prototype plasma (Gabor) lens was built at Imperial College London. The goal was that it should operate at high electron densities. A detailed description of the configuration of the lens and the experimental investigations of the characteristics of the plasma that the prototype Gabor lens contained can be found elsewhere [54, 222]. Nonetheless, an outline of the main components of the lens is included here.

A schematic diagram of the prototype Gabor lens is shown in fig. 4.1. The total length of the lens was 540 mm from end flange to end flange. The longitudinal confinement of the plasma was provided by a pair of grounded cathodes at the two ends of the lens and by a central anode connected to a source of positive high-voltage. The three electrodes were formed out of cylindrical copper tubes. The central anode had a length of 444 mm and an inner diameter of 85.7 mm. Four rows of four 10 mm-diameter holes were made along four lines parallel to the axis of the anode and spaced by $90°$ so that the vacuum inside the anode would be maintained effectively. The anode was electrically isolated from the vacuum tube of the lens by two ceramic spacers which also maintained the position of the central electrode. A high-voltage copper connector was soldered to the anode to provide a feed-through designed for plasma-confining voltages up to 60–65 kV.

The two end electrodes had a diameter of 66.7 mm, smaller than that of the anode, and a length of 34 mm. The minimum gap between any of the two end electrodes and the central anode was 16.8 mm. To further reduce the likelihood of sparking, the ends of the grounded electrodes were manufactured with rounded edges. The end electrodes were press fitted into the mild-steel end flanges. An outer tube covered the vacuum tube of the lens and was electrically connected to ground together with the end flanges.

The radial confinement of the plasma inside the lens was ensured by an axial magnetic field generated by a series of pancake coils wound around the vacuum tube. The configuration of the pancake coil is shown in fig. 4.1. The water-cooled copper conductor with a square cross-section was wound in four principal layers. The number of windings was locally increased to seven at specific locations to improve the uniformity of the magnetic field. The coil produced a maximum magnetic flux density of approximately 55 mT for a current of 45 A. More typical values of the current through the coil were in the range 14–30 A. The high-voltage supply
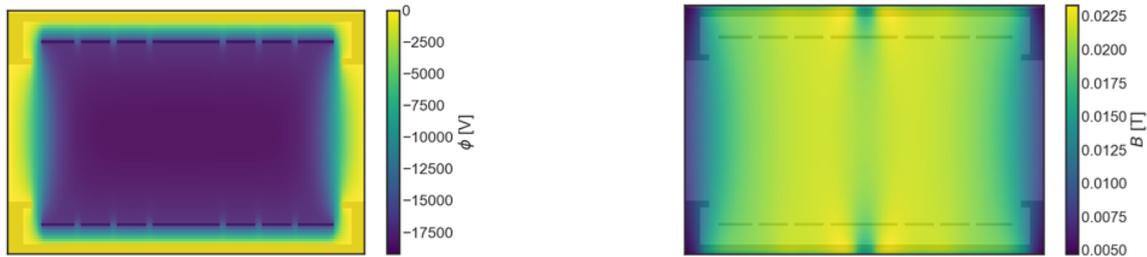
Figure 4.2: Electrostatic potential map (left) and magnetic flux density map (right) in the longitudinal mid-plane of the IC lens prototype as extracted from VSim and FEMM. The central anode, end electrodes, and the vacuum tube are overlaid. The anode potential, $17\,\mathrm{kV}$ is matched to the magentic flux density in the centre of the lens, $20.4\,\mathrm{mT}$ (coil current $20\,\mathrm{A}$).

provided typical values between $8\,\mathrm{kV}$ and $20\,\mathrm{kV}$ to the central anode. The vacuum pressure inside the lens was of the order of $1 \times 10^{-5}\,\mathrm{mbar}$ when the plasma was established, while a minimum of $3 \times 10^{-7}\,\mathrm{mbar}$ was achieved in the absence of plasma.

An example of the electrostatic potential generated inside the lens by the cylindrical electrodes is shown in fig. 4.2. Due to the short separation of the anode and the end electrodes, the equipotential lines are curved and reach into the internal volume of the anode. Thus, a high-density plasma with the diameter comparable to that of the anode is expected to have curved edges at the two ends of the lens. The circular perforations in the wall of the anode lead to small regions where the potential is lower than the average value in most of the internal volume of the central electrode.

Figure 4.2 also shows the magnitude of the magnetic flux density in the mid-plane of the lens obtained with a finite-element calculation in FEMM [223] and a model of the pancake coils of the lens. The spatial variation of the number of layers of the coil leads to regions of non-uniform magnetic field. The magnitude of the field is lower at the centre of the lens due to the absence of a current flowing around the anode in the mid-plane of the lens, as found in simulation and shown in fig. 4.2. The non-uniform field strength is a consequence of the fact that the coil is interrupted near the mid-plane to create the space required for the high-voltage connector to reach the anode from outside of the vacuum tube and the lens. Thus, the lower magnetic field around the mid-plane of the lens can lead to stronger radial transport and higher electron loss in the middle of the lens compared to the rest of its active volume.

The configuration of the lens shown in fig. 4.1 differs from the one initially proposed by Gabor [49] through the absence of a thermionic cathode, or any other type of source of electrons. Thus, the plasma was produced in the lens through an uncontrolled, spontaneous electric discharge in the non-ideal vacuum of the residual background gas. The typical procedure to turn on the lens was to increase the high voltage applied to the anode and the current through the coil such that a plasma would be generated from a spontaneous discharge event. The establishment of a high-density plasma was typically observable by a significant increase in the vacuum pressure of the lens. Furthermore, the formation of an electron plasma was validated by measuring the current of the ions that leave the lens on axis once a negative space-charge was accumulated inside the lens [54]. Several measurements carried out on distinct days [54] found a consistent variation of the ion current that leaves the

lens with the current through the coil. Thus, it was concluded that the plasma had similar characteristics every time it was produced by turning the lens on following the procedure described above.

In addition to the tests of the prototype carried out in a laboratory at Imperial College London, the lens was tested with $1.4\,\mathrm{MeV}$ proton beams at the Ion Beam Facility at the University of Surrey [224]. The prototype was found to convert thin pencil-beams into rings which suggested the generation of an unstable plasma inside the lens. The work presented in this chapter is part of a campaign of particle-in-cell (PIC) simulations aimed at understanding the characteristics of the plasma that could lead to the focusing effects observed experimentally and the cause of plasma instability. The main results of the beam test are summarised in section 4.3 together with the further insight that was obtained from PIC simulations of the lens.

The particle-in-cell (PIC) method is a powerful technique to conduct simulations of weakly interacting particles and to study the realistic behaviour of a plasma by modelling the motion of significantly fewer particles than are present in the physical system. For the interested reader, a detailed description of the essential parts of a PIC simulation is given in appendix F. The main principles of the PIC procedure are summarised in the following. The method employs finite-sized macro-particles that interact via the correct Coulomb potential at long ranges beyond the overlap distance and via a reduced interaction at small distances to correct for the effect of using fewer particles. A macro-particle represents a large number of particles that are near each other in phase space. Within the PIC method, each macro-particle is assigned a shape function that defines how charge density is distributed about its centre. Furthermore, the simulation volume is divided into cells by a grid to which the values of the electromagnetic fields are attached. The shape function ensures the correct interpolation of the field values from the grid to the location of the macro-particles and, similarly, the correct deposition of the charge from the continuum onto the grid nodes. The fields and the macro-particles are advanced in time based on discretisation of the equation of motion and Poisson's equation in a computation cycle that is repeated at each time step. A more detailed illustration of the link between the mathematical model of a plasma and its numerical solution can be found in [225].

Starting from the capabilities of the PIC method, a series of systematic simulations of electron plasmas confined by the prototype Gabor lens was performed which is presented in the following two sections.

4.2   Stable-plasma regime

In order to evaluate the capability of the prototype lens to produce a highly uniform electron plasma, and, thus, a linear focusing force, a PIC simulation was carried out of the electrons in the lens. It was assumed that the plasma discharge led to the entire volume of the lens to be filled with electrons. The simulation was aimed at the evolution of the plasma after its production. The geometry of the prototype lens was reproduced in VSim according to the engineering drawings. The model included the central anode, the two end electrodes, and the vacuum tube with the end flanges. The width of the walls of the anode and end electrodes was set to double that

Table 4.1: Main parameters of the simulation of the prototype plasma lens in a stable regime.

Parameter	Value
Current in the coil	20 A
Average magnetic field (B_{avg})	20.4 mT
Anode potential (V_A)	17 kV
End electrodes potential	0 kV
Maximum theoretical electron density ($n_{e,\max}$)	2×10^{15} m^{-3}
Initial filling factor ($n_e/n_{e,\max}$)	10%
Number of time steps per plasma period	30
Grid size	12 cm $\times$ 12 cm $\times$ 54 cm
Number of cells in each direction	$200 \times 200 \times 160$
Macro-particles per cell	10
Background gas (H_2O) pressure	5×10^{-6} mbar

from the drawings to ensure a proper meshing of these surfaces.

To model the lens accordingly, the configuration of the coil together with all the metallic surfaces were replicated separately by using a finite element analysis package, FEMM [223]. The software is used to solve 3-D axisymmetric linear magnetostatic problems with a low computational cost. For the studies presented here, the package was used to obtain a 2-D cylindrically symmetric magnetic field map starting from the value of the current through the coil (see fig. 4.2). The field map was then loaded into VSim as the longitudinally confining magnetic field with the help of a regular grid interpolator and the interface between Python and VSim.

The main plasma parameters and lens specifications used in the simulations are given in table 4.1. The anode potential was matched to the average magnetic field according to the work function of the lens [226]

$$V_A = \frac{eR_A^2}{8m_e}B^2 \tag{4.1}$$

H_2O was assumed to be the most probable background residual gas. It was introduced in the simulation as a neutral fluid with three interactions with the surrounding electrons modelled during the run: elastic collision, impact excitation, and impact ionisation. The corresponding cross-sections were obtained from [227] and [228]. Furthermore, secondary electron emission was included from the surface of the anode and end electrodes. Both the electrons and the ions produced in the impact ionisation could be absorbed at any of the metallic surfaces, as well as at the entry and exit plane of the lens inside the beam-pipe. A typical simulation required a running time of approximately 12 hours on 20 cores to track the evolution of the plasma for a time interval of 10 μs. VSim was run on the high-performance computing cluster SCARF [229] hosted at STFC's Rutherford Appleton Laboratory.

The simulation volume was bounded by Dirichlet boundary conditions with a voltage set to 0 V. The vacuum tube was assumed to fill the full space between its inner cylindrical surface and the planes that enclose the simulation volume. A perfect electric conductor (PEC) material with Dirichlet boundary conditions was used for the anode and the two end tubes, thus allowing the voltage on the electrodes to be fixed by the user.

The initial distribution of the electron cloud was set out starting from two different assumptions:
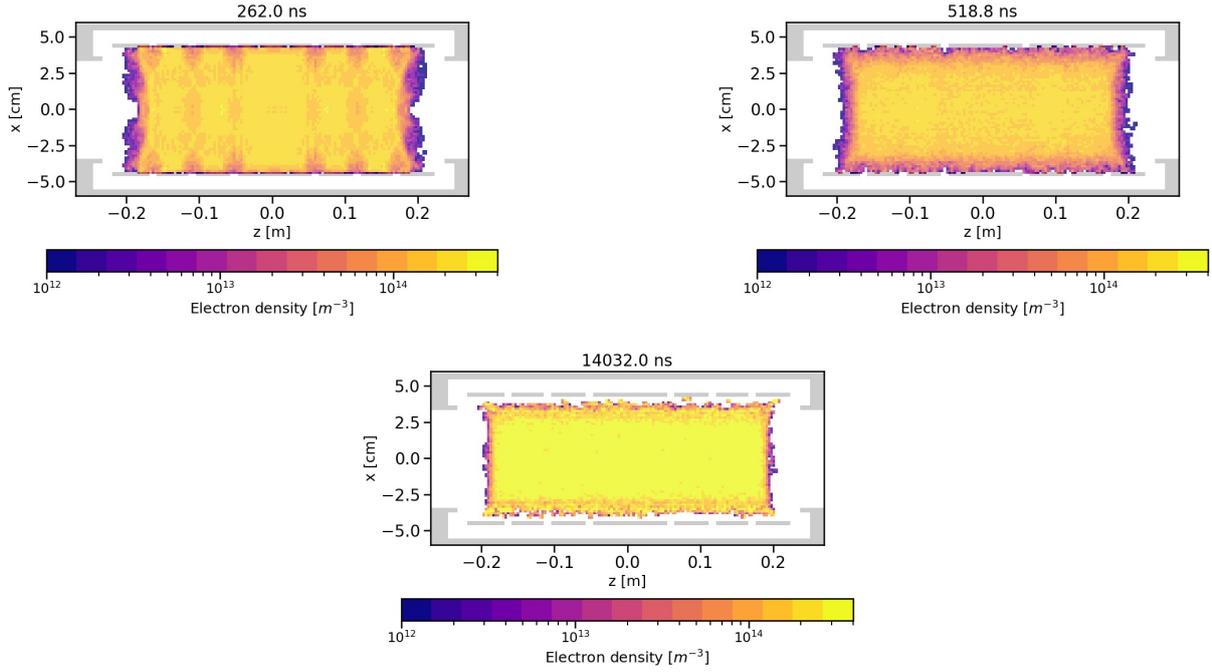
Figure 4.3: Electron-plasma density in the longitudinal mid-plane of the IC lens prototype extracted at two intermediate (left and middle) and the last time step (right) of a PIC simulation in VSim. Low-density regions appear close to the circular perforations in the central anode of the lens. The grey areas indicate the conducting surfaces of the lens electrodes.

(i) The first method was to assume that the plasma starts from thermal equilibrium. A temperature in the range $10\,\mathrm{eV}$ to $100\,\mathrm{eV}$ was chosen to match previous observations and simulations [230]. The velocity components for each electron were considered to be independent and were sampled from a Maxwellian distribution. The electron density was assumed to be constant.

If the plasma initially occupies only a small volume near the middle plane of the lens, the electrons were seen to flow towards the two ends of the lens until an 'equilibrium length' was reached as exemplified in the bottom row of fig. 4.3. If the entire inner volume of the anode is filled with electrons at the initial time step of the simulation, two plasma waves propagate from the edges of the plasma column towards the centre. The waves continued to propagate forwards and backwards for the entire duration of the simulation as regions with higher plasma density. To avoid the formation of these plasma waves, the electrons had to be initialised such that the length of the plasma column at $t = 0$ is close to its 'equilibrium length'.

(ii) For the second approach, the electron cloud is assumed to start from one possible equilibrium state. A particular class of equilibrium is chosen for which the plasma rotates around its axis of symmetry as a rigid column with angular frequency [190]

$$\omega_E^- = \frac{\omega_{ce}}{2}\left[1 - \left(1 - \frac{2\omega_{pe}^2}{\omega_{ce}^2}\right)^{\frac{1}{2}}\right] \tag{4.2}$$

where ω_{pe} is the electron plasma frequency, ω_{ce} is the electron cyclotron frequency, and ω_E^- is character-
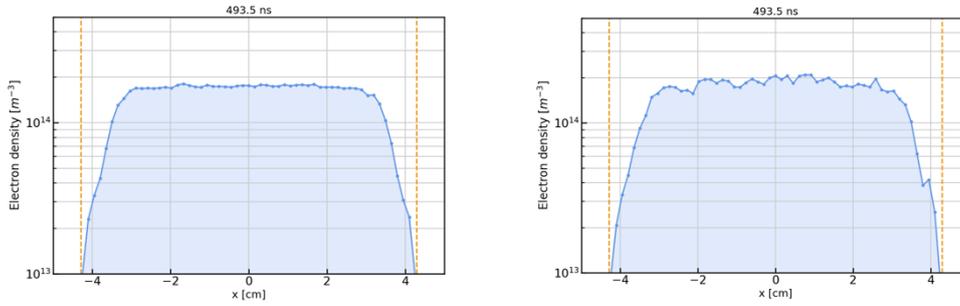
Figure 4.4: Comparison of the transverse electron density profile in the mid-plane of the lens after approximately 500 ns for two different initial plasma distributions: thermal equilibrium (left) and rigid rotor motion (right).
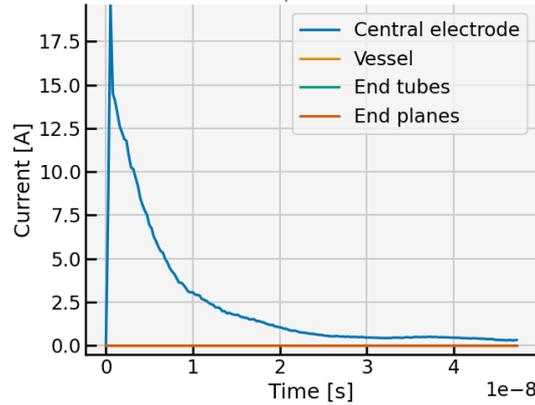
Figure 4.5: The absorbed electron current at the surface of the components of the lens as a function of time.

istic to an $\mathbf{E} \times \mathbf{B}$ rotation of the plasma column.

Figure 4.4 shows a comparison of the radial electron density profile after approximately 500 ns for the two methods of initialising the non-neutral plasma in the simulation. Both initial conditions led to similar main features: pseudo-flat plateau of the electron density below $2 \times 10^{14}\,\text{m}^{-3}$ with a radius close to 3 cm, and a negative radial gradient near the inner surface of the anode. Initial thermal equilibrium leads to a more uniform density plateau, while the initial rigid rotor state results in a more non-uniform density distribution on the time scale considered here. For the rest of the simulation, the plasma column slowly continues to contract radially which is a suggestion that the state of the plasma is not in full equilibrium and there is a residual transient behaviour most probably dictated by the initial distribution of the plasma column. However, the plasma was observed to evolve towards a more stable state. The transient behaviour is most strongly visible in the number of electrons lost from the lens immediately after the start of the simulation while the plasma redistributes itself. Figure 4.5 shows a high spike in the absorbed electron current on the inner surface of the central electrode during the first 20 ns of the simulation.

Figure 4.3 shows a typical electron cloud under a quasi-stable regime after more than 10 µs. The electron cloud has a radial extension smaller than the anode radius. A pseudo-flat plateau is observed at $\sim 60\%$ of the nominal electron density between ± 2 cm. The negative radial gradient of the electron density is consistent with a stable plasma. In the transverse plane, the plasma is rotating azimuthally around the central axis of the lens. The electrons are lost from the two ends of the plasma column and near the inner surface of the anode
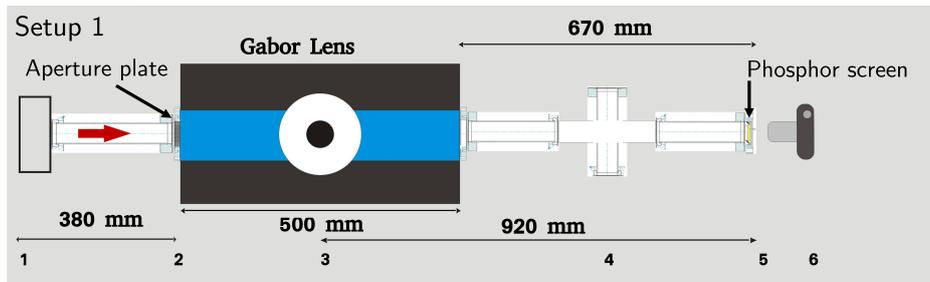
Figure 4.6: Schematics of the main setup used during the beam test of the Gabor lens prototype. The drawing includes the lens, aperture plate, beam pipes, and the imaging detector. Figure reproduced from [54].

in approximately the first $20\,\mathrm{ns}$. No electrons were registered to have reached the end planes of the lens or the walls of the outer vacuum vessel. Moreover, fig. 4.3 shows that, in the initial part of the simulation, while the electrons redistribute in the lens, some low density regions appear near the holes of the central electrode. As the electrostatic field lines leak through these perforations, some high-energy electrons can escape outside of the electrode. The low density regions near the holes disappear as the plasma column contracts radially.

The PIC simulations presented above indicate that the prototype Gabor lens was able to confine a highly uniform electron plasma with densities in the vicinity of 40% of the Brillouin limit provided that the mechanism by which the plasma was formed filled the entire volume of the lens with a highly uniform electron cloud. The uniform plateau of the plasma density, and, thus, the region of linear focusing extends beyond a radius of $3\,\mathrm{cm}$. The main region of electron loss was identified as the inner surface of the anode and the circular perforations in its wall. Nevertheless, the linear focusing force produced by the stable plasma predicted by the simulations was not recorded in practice in the experimental tests of the prototype lens. A separate set of simulations was carried out to explore the more complex plasma dynamics inside the lens.

4.3   Unstable-plasma regime

4.3.1   Beam test

The focusing capabilities of the prototype Gabor lens were investigated at the Ion Beam Facility at the University of Surrey [224] in a beam test using protons with kinetic energy of $1.4\,\mathrm{MeV}$. A schematic diagram of the experimental setup is shown in fig. 4.6. The main proton beam was converted into several thin pencil-beams with the use of an aperture plate positioned at the entrance of the lens. The plate contained $2\,\mathrm{mm}$-diameter holes arranged in a spiral pattern designed to minimise the overlap of the outgoing pencil-beams under a focusing force that is rotationally symmetric about the beam axis.

Once the main beam was converted into pencil-beams, the protons passed through the lens during multiple measurements either in the presence of an electron plasma in the lens or with the external fields of the lens turned off. After a further section of evacuated beam-pipe, a phosphor screen was installed downstream of the lens as indicated in fig. 4.6. The light produced by the protons hitting the screen was recorded with a DSLR camera using an exposure long compared to the duration of a beam spill.
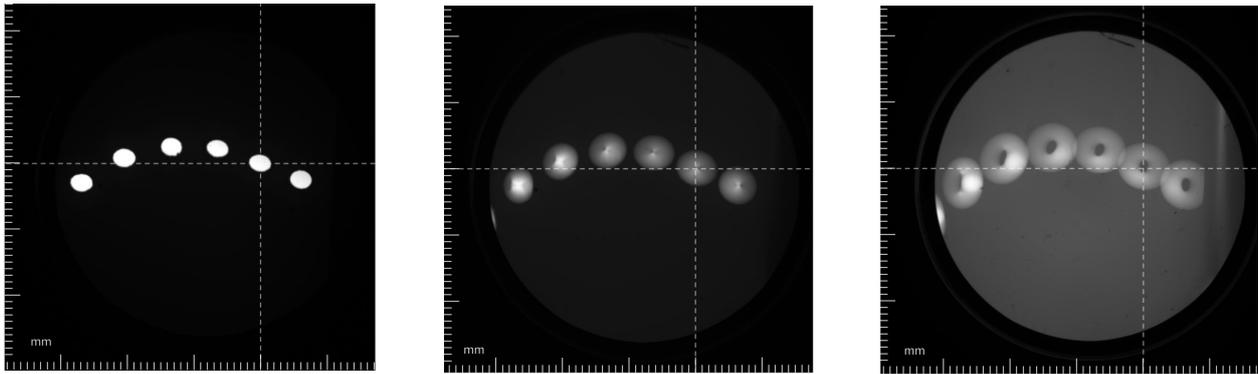
Figure 4.7: Camera images of the six beam spots beyond the aperture as recorded during the beam test with the lens off (left) and on at a current through the coils of $28\,\mathrm{A}$ (middle) and $33\,\mathrm{A}$ (right). A voltage of $20\,\mathrm{kV}$ was applied to the anode when the lens was on. The dashed lines indicate the beam axis and the central beam spot. Figure reproduced from [54].

In the absence of plasma (coil current and anode high-voltage turned off), the proton pencil-beams propagate through a drift region from the aperture plate to the screen. Six pencil-beams were imaged with the lens turned off as shown in first photograph of fig. 4.7. The axis of symmetry of the lens passes through the centre of the second pencil-beam from the right. A comparison of the shape and diameter of the beam spots on the screen with the circular apertures of the plate allowed the divergence of the beam to be determined. Thus, the divergence was calculated to be $1.6\,\mathrm{mrad}$ and $0.5\,\mathrm{mrad}$ in the x and y directions, respectively. The position of the spots on the screen were traced back to the entrance of the lens and used in combination with the beam divergence in the particle-tracking simulations presented later in this chapter.

Separate images of the screen were recorded with the external fields of the lens turned on. Two images are shown in fig. 4.7 for the lens operating at $20\,\mathrm{kV}$ and two coil currents of $28\,\mathrm{A}$ and $33\,\mathrm{A}$. The presence of plasma inside the lens can be inferred from the differences in the two images for different coil currents. A higher current in the pancake coils corresponds to a stronger average magnetic field and, thus, a denser plasma can be confined by the lens, leading to a stronger focusing of the pencil-beams. However, fig. 4.7 shows the main anomalous focusing effect that was observed repeatedly during the beam test of the prototype Gabor lens. In short, the proton pencil-beams were converted into ring-like structures on the phosphor screen. The brightness of the images was found to vary around the circumference of each ring. More details on the measurements performed during the beam test to characterise the observed rings can be found in [54]. Similar behaviour was reported in a previous test of the lens with a single proton beam [222]. These experimental observations led to a series of simulations of the prototype Gabor lens and the plasma generated by it aimed at identifying the cause of the anomalous focusing effect on the proton beams. The simulations are described in the rest of this chapter.

4.3.2   Simulation of protons passing through an unstable plasma

The conversion of the circular pencil-beams into rings by the prototype lens led to the hypothesis that the lens generated an unstable plasma during the beam tests. As the rings were observed repeatedly, even on

consecutive days between which the lens was turned off and on, a similar plasma instability had to be generated and maintained inside the lens each time the lens was operated. Based on the long duration of the instability in the plasma, the diocotron mode [217] was proposed as the cause of the experimental observations. The diocotron instability is one of the most ubiquitous instabilities [231] that can appear in a long non-neutral plasma column and has been extensively studied in Penning-Malmberg traps. The diocotron modes refer to azimuthal electrostatic flute perturbations (longitudinal wave number equal to zero) in the potential of the plasma, and, as a result, to azimuthally rotating distortions in the density profile of the plasma. The fundamental diocotron mode (lowest mode number) can also be associated with the bulk rotation of an off-axis plasma column that rotates around the axis of the trap because of the interaction between the plasma and its image charge induced in the conducting walls of the electrodes.

The diocotron modes are particularly distinguishable because of the fact that they are *negative energy modes* [187]. The distorted plasma has lower potential energy compared to the cylindrically symmetric state of equilibrium of the plasma with the same density, and has negligible excess kinetic energy. Therefore, the total energy of the plasma during the diocotron instability is lower than that the energy of the same plasma in an equilibrium state. As a consequence, the diocotron modes can be driven unstable by processes that remove energy and angular momentum from the plasma, such as collisions with the residual background gas or the finite resistance of the walls of the electrodes [187].

A second important feature of the diocotron modes is that they are driven by specific radial profiles of the plasma density. In particular, hollow plasma profiles and negative radial gradients in the density [232] are well-known to make the plasma prone to this instability. Accordingly, the simulations of the lens prototype that are presented below were set-up with initial distributions of the plasma density that are known to lead to diocotron-type instabilities. A few different initial distributions were investigated in relation to the evolution of the plasma that was observed in simulation. Moreover, thin proton pencil-beams were propagated through the unstable plasmas to arrive at a correlation between the plasma instabilities of interest and their focusing effect on proton beams.

The time evolution of unstable non-neutral electron plasmas inside the prototype Gabor lens was studied with the particle-in-cell simulation method and the VSim software [233]. The geometry of the central anode, the two end electrodes, and the vacuum tube of the lens was reproduced in VSim. The axial magnetic field was treated as a static field and was imported at the start of the simulation from a 2-D cylindrically symmetric map calculated with a separate finite-element-analysis package. The voltage on the central anode was chosen such that the maximum electron density from the longitudinal confinement condition to be equal to the maximum electron density imposed by the condition for radial confinement (see eq. (3.22)). In other words, the magnitude of the anode voltage was matched to the value of the magnetic field at the centre of the lens to ensure that both the radial and longitudinal confinement conditions lead to the same maximum plasma density. The time step was set to $0.2\,\mathrm{ns}$ (approximately $6\tau_c$, where τ_c is the cyclotron period). The grid had a transverse cell size $\Delta x = 0.14\lambda_D$ and longitudinal cell size $\Delta z = 0.8\lambda_D$, where λ_D is the Debye length of the plasma. Each
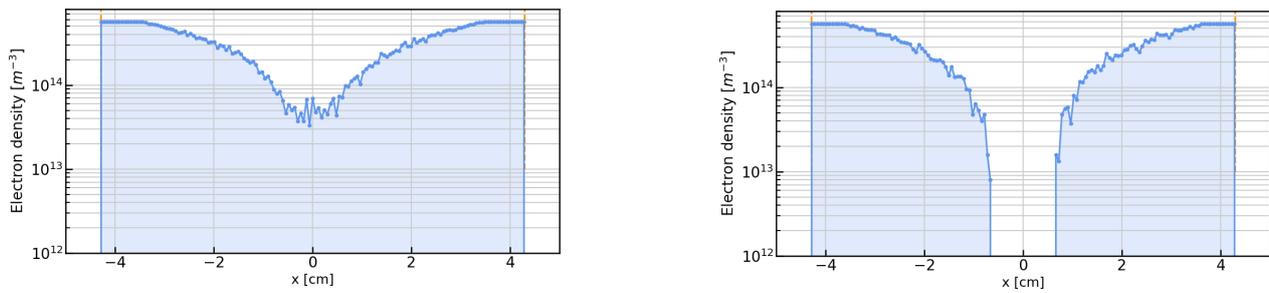
Figure 4.8: Initial radial profile of the plasma density that was observed in PIC simulations to evolve into a ring- (left) and a diocotron- (right) type instability. The inner surface of the cylindrical anode is situated at a radius of $4.3\,\mathrm{cm}$.

cell was initialised with 64 macro-particles. Peak plasma densities in excess of $1 \times 10^{14}\,\mathrm{m}^{-3}$ were considered here based on some previous estimations of the plasma density that was produced during the beam test of the prototype lens. The evolution of the plasma was typically simulated until about $2\,\mathrm{\mu s}$. By contrast, the transit time of a proton with a kinetic energy of $1.4\,\mathrm{MeV}$ through the active volume of the lens is approximately $27\,\mathrm{ns}$. The proton beam was considered to be continuous (unbunched).

In addition to the electron macro-particles, protons were also added to the simulations as thin pencil-beams. Eight circular beams were generated at the entrance of the lens according to the spatial pattern and divergence calculated from the measurements done during the beam test. The initial position of the beams can be seen in fig. 4.9 and 4.12. The second beam from the left is aligned with the axis of symmetry of the lens. Once the source of the proton macro-particles was turned on, the beams were continuously generated until the end of the simulation. The delay between the start of the plasma simulation and the start of the generation of the beams at the entrance of the lens could be varied by the user. Throughout the full duration of a simulation, the position and velocity of each proton macro-particle that hit the exit plane of the lens was recorded to a file. The resulting proton phase-space distribution was then used to track the protons with BDSIM [141] for a further drift space of $67\,\mathrm{cm}$ until a plane situated at the same position as the phosphor screen in the setup of the beam test. The time-integrated transverse distribution of the protons on this plane was finally written to file and plotted.

Hollow electron ring

The first unstable plasma mode explored in simulation was an electron column with a lower particle density along the axis of symmetry compared to the edge of the plasma. The initial radial profile of the electron density is shown in fig. 4.8. The density around the axis is more than an order of magnitude lower than the peak magnitude at the edges. The negative radial gradient leads to perturbations in the electron density which propagates azimuthally. In the first $100\,\mathrm{ns}$ of the simulation, the plasma becomes more compact radially and the electrons move away from the conducting surface of the anode. However, the initial radial gradient is not large enough to drive strong azimuthal plasma waves and the electron column retains its ring-like structure for the full duration of the simulation. The density of the electron column in the mid-plane of the lens is shown at a few representative time steps in fig. 4.9. Plasma 'lobes' and azimuthal waves are noticeable at specific time steps,
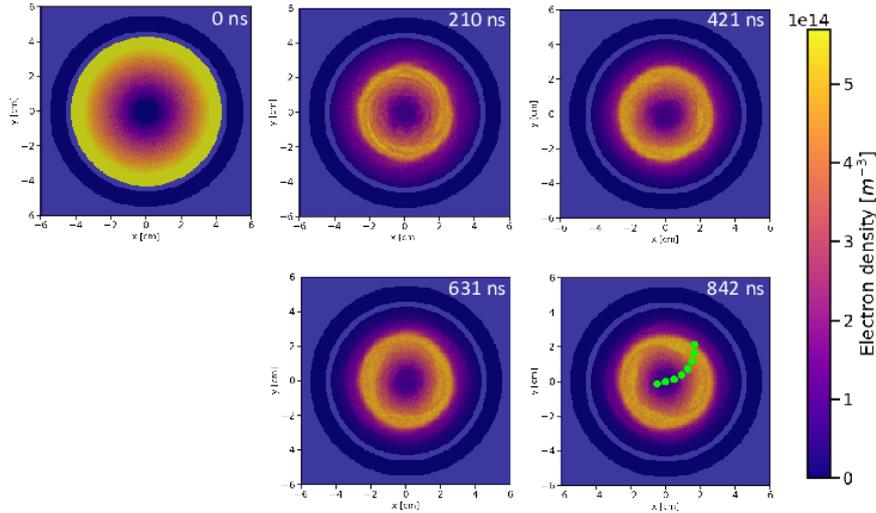
Figure 4.9: The time evolution of an electron-plasma instability caused by an initial density distribution with a positive radial gradient as simulated with VSim. Each frame is a transverse cross-section through the mid-plane of the lens at the time indicated in the top-right corner. The green spots mark the entry position of thin proton pencil beams.

but they decay rapidly. As a result, the focusing force of the electrons on the proton beams is highly cylindrically symmetric about the axis of the lens.

Figure 4.11 shows the density of the proton macro-particles that hit the virtual screen downstream of the lens after passing through the ring-shape plasma described in the previous paragraph. The trajectories of the protons that propagate inside the region of low density and close to the axis are not influenced by the plasma. By contrast, the pencil-beams that travel further away from the axis are distorted by the electrons. These beams create elliptical bright spots on the screen and show no sign of smearing. In order to generate stronger diocotron plasma waves, the radial gradient of the initial electron density was increased further by removing the electrons from the central region of the plasma column.

Diocotron instability

The second unstable plasma that was explored in a simulation was generated by a larger radial gradient of the initial plasma density and an initially completely hollow electron ring, as showed in fig. 4.8. By approximately 500 ns, the plasma perturbations take the shape of azimuthal lobes that rotate around the axis of the lens without the full fragmentation of the plasma ring that is commonly observed in experiments (see, for example, [234]). The plasma density in the mid-plane of the lens is shown in fig. 4.12 at several representative time steps. The approximate three- and four-fold symmetry observed throughout the simulation suggests that the third and fourth azimuthal diocotron modes were predominately excited. As these modes break the cylindrical symmetry of the focusing force in the lens, they could lead to smearing of the proton beams that pass through the unstable plasma. Furthermore, the density perturbations rotate in the simulation with periods of several hundred of nanoseconds
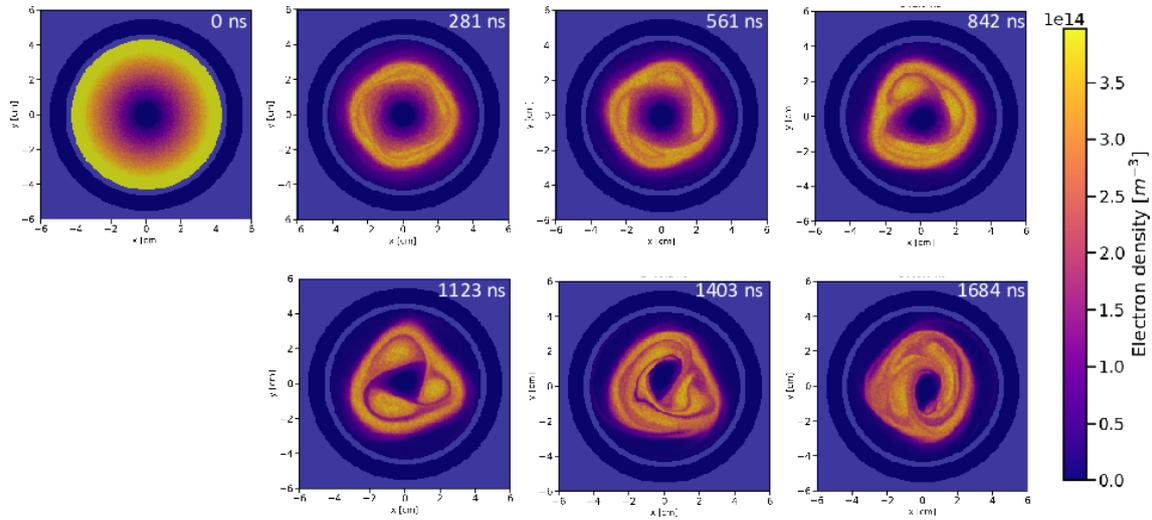
Figure 4.10: The time evolution of an electron-plasma instability caused by an initial density distribution with a large positive radial gradient as simulated with VSim. Each frame is a transverse cross-section through the mid-plane of the lens at the time indicated in the top-right corner. The azimuthal shear in the plasma and the formation of 'lobes' are indicative of the diocotron instability.
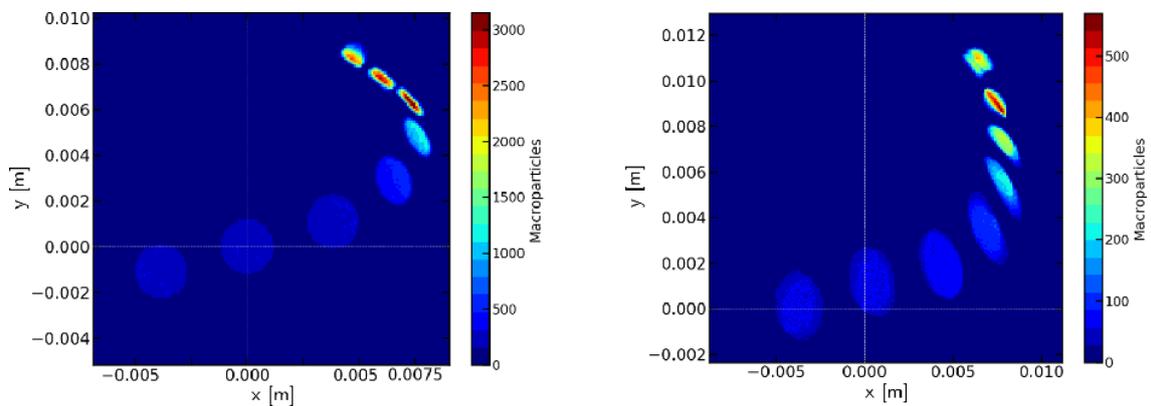
Figure 4.11: Transverse profiles of eight initially circular proton pencil-beams tracked through the ring (left) and diocotron (right) instabilities shown in fig. 4.10 and 4.9, respectively.

which is significantly larger than the transit time of the protons. Thus, the rotation of the high-density regions of the plasma may not average out the asymmetric focusing effect.

To test the above hypothesis, proton pencil-beams were generated at the entrance of the lens from the start to the end of the simulation described in the previous paragraph. The positions of the protons that reached the exit plane of the lens were recorded and the time-integrated distribution at the location of the virtual phosphor screen is shown in fig. 4.11. Compared to the case of the hollow plasma in fig. 4.9, the stronger diocotron waves shown in fig. 4.12 generate an asymmetric focusing force on the pencil-beams that propagate close to the axis of the lens. The corresponding protons are smeared on the screen creating elliptical bright spots. Similar smearing is observed for the pencil-beams that travel further away from the axis of the lens. These beams also experience stronger focusing as the bright spots on the screen are smaller compared to the ones created by the central beam. Nonetheless, the formation of rings could not be reproduced with an unstable plasma that evolved similarly to a diocotron mode with an azimuthal mode number $m_\theta > 1$. After approximately $2\,\mu$s, the plasma instability grows rapidly and a large fraction of the electrons are lost on the inner surface of the anode.

Given the results of the simulations presented so far in this section, it can be argued that a plasma distribution with an even stronger asymmetric focusing force is needed for the pencil-beams to be converted into rings. One such plasma distribution is the fundamental, $m_\theta = 0$, diocotron mode for which the entire electron column is radially displaced from the central axis of the lens. A simple model for the fundamental diocotron mode is presented in sub-section 4.3.3. Before describing the model in full, a third and final example of an unstable plasma simulated with the particle-in-cell method is reported below. The time evolution of the plasma in this last example is highly similar to the $m_\theta = 1$ diocotron mode.

'Dipole' instability

The third example of an unstable plasma was driven in a simulation by the electron density shown in the top left frame of fig. 4.12. The aim of the initial profile of the electron density was to create an average beam-focusing force that is not cylindrically symmetric, while keeping the negative radial gradient that typically leads to instabilities. In a realistic scenario, such a plasma distribution can be created by the axial loss of electrons at the two ends of the lens from a region that does not coincide with the central axis. Alternatively, a similar plasma density can be established by a surplus of electrons created close to the wall of the anode by cascade ionisation of the residual background gas or by secondary emission from the metallic surface of the anode.

The rest of the frames in fig. 4.12 show the evolution of the instability which was labelled as a 'dipole' instability due to the two regions of high and low electron density. The rotation of the two regions around the axis of the lens is similar to the rotation of a uniform plasma column that is displaced from the central axis (i.e. the $m_\theta = 1$ diocotron mode). As the plasma rotates about the beam axis, the region with the highest density gradually moves radially towards the centre of the lens. After approximately $2\,\mu$s the plasma looses its initial dipole structure and changes into a cylindrically symmetric electron column which remains stable.

The displacement and rotation of the bulk of the plasma observed in fig. 4.12 leads to an equivalent rotation
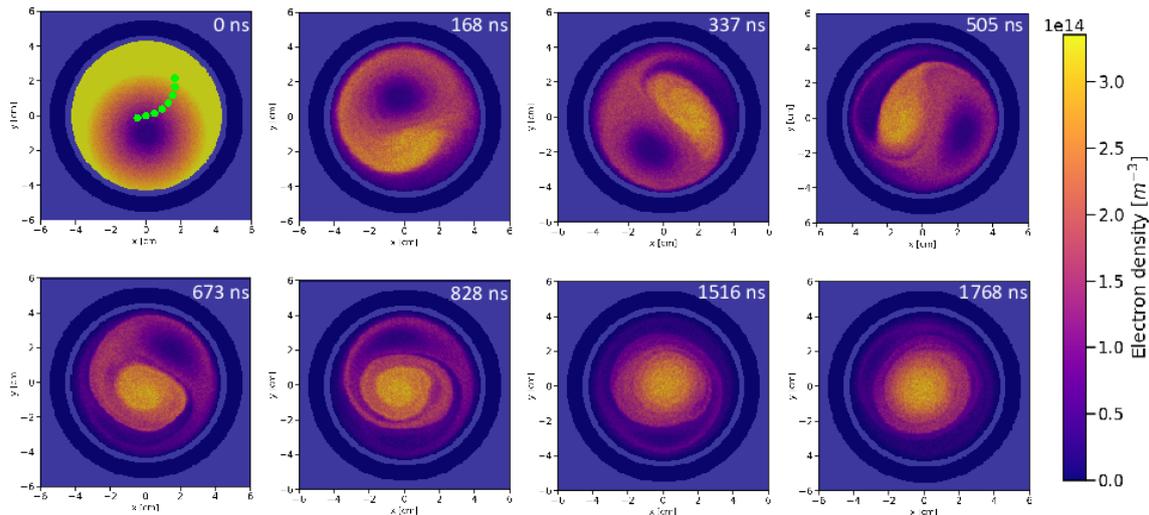
Figure 4.12: The time evolution of an electron-plasma instability caused by an asymmetric initial density distribution with a positive radial gradient as simulated with VSim. Each frame is a transverse cross-section through the mid-plane of the lens at the time indicated in the top-right corner. The green spots mark the entry position of thin proton pencil beams.

of the focusing centre of the lens. For example, protons entering the lens on axis at approximately $170\,\text{ns}$ after the start of the instability exit the lens with a negative vertical displacement from the central axis. By contrast, the protons passing through the lens with an approximate delay of $350\,\text{ns}$ experience an average focusing force directed in the positive vertical direction. The result of the rotation of the focusing centre of the lens is further validated in fig. 4.13.

As before, eight proton pencil-beams were tracked from the entrance to the exit of the lens and, then, downstream to the location of a screen. The protons were injected into the unstable plasma either from the beginning of the instability shown in fig. 4.12 or with a time delay of $1\,\mu\text{s}$. The unstable plasma led to bright spots on the screen significantly larger in diameter compared to the initial pencil-beams. The bright spots appear as a consequence of the screen being 'painted' by the pencil-beams as the focusing centre of the lens rotates simultaneously with the rotation of the high-density bulk of the plasma around the axis of the lens.

When the protons were injected from the start of the simulation, the pencil-beams generated ring-like bright regions on the screen as shown in the plot on the right of fig. 4.13. The pencil-beams near the axis of the lens have the largest separation with respect to the centre of mass of the plasma and, thus, experience the largest average focusing force. As a result, the radii of the rings that are produced by the beams closer to the axis of the lens are larger compared to those of the rings that sit further away from the axis. More than eight rings appear overlapped on the virtual screen since the focusing centre of the plasma slowly approaches the axis of the lens as the instability is gradually damped. The centre and radius of each ring slowly varies in time leading to the pattern in fig. 4.13.

In a separate run, the protons were injected into the lens with a time delay of $1\,\mu\text{s}$ to verify that the rings
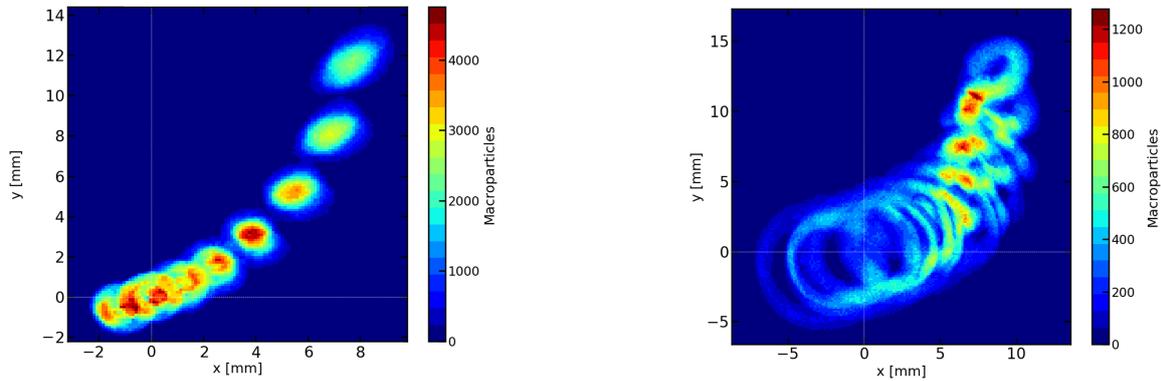
Figure 4.13: Number of proton macro-particles hitting a screen $67\,\mathrm{cm}$ downstream of the lens. The eight initially circular pencil beams passed through the 'dipole instability' shown in fig. 4.12. The protons entered the lens at a time $\Delta t = 1\,\mathrm{\mu s}$ (left) and $\Delta t = 0\,\mathrm{\mu s}$ (right) after the start of the PIC simulation of the electron plasma.

appear as a result of the bulk of the plasma rotating around the axis of the lens in the first half of the simulation. Under these conditions, the plot on the left of fig. 4.13 shows that the rings were not produced on the screen anymore.

In conclusion, a few distinct unstable plasmas were observed in particle-in-cell simulations of the prototype Gabor lens built at Imperial College London. The instabilities were generated by suitably chosen initial distributions of the electron density that exhibit either large negative radial gradients or a radial displacement of the bulk of the plasma from the axis of the lens. The focusing effect of these unstable plasmas on thin pencil-beams was assessed by injecting proton beams into the lens as part of the PIC simulations simultaneously with the time evolution of the plasma. It was found that the pencil-beams are converted into rings only in the case for which the bulk of the plasma is displaced from the axis of the lens and rotates around it. The following sub-section introduces a simplified analytic model of the electric field produced by a plasma column that rotates around the axis of a Gabor lens.

4.3.3 Model of a rotating off-axis plasma column

The particle-in-cell simulations presented above indicated that the $m_\theta = 1$ diocotron plasma instability could be the source of the anomalous focusing effect that was observed during the beam test of the prototype Gabor lens. However, it was desirable to obtain a better agreement between the predictions from the simulations and the images recorded during the beam test in order to gain more insight into the capabilities of the lens, for example, regarding the range of plasma densities achieved during the operation of the lens. An attempt to better reproduce the results of the beam test in simulations required parameter scans of the plasma density, radial offset, and rotation frequency of the electron column. Since a single PIC simulation of an unstable plasma such as the one presented above typically runs on 16 CPU cores for up to 48 hours, the PIC method is too slow for such parameter scans. Therefore, a simplified analytical model is presented here that allowed faster particle-tracking simulations of protons passing through a rotating plasma column.
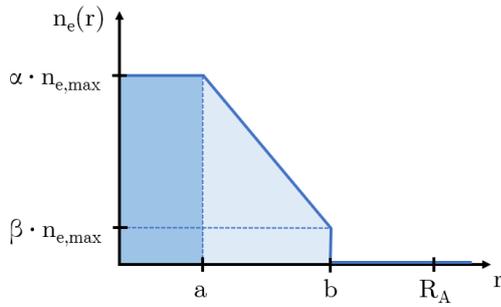
(a) Negative radial gradient in the electron density.
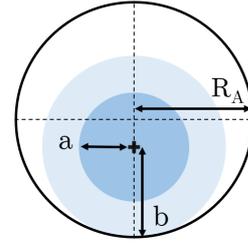
(b) Transverse view of the plasma column inside the anode of the lens.

Figure 4.14: Radial electron density profile for a simplified model of a long plasma column displaced from the symmetry axis of the lens. R_A represents the radius of the lens anode and $n_{e,max}$ is the theoretical maximum electron density. r is measured from the centre of the plasma column. The profile is fully described by the parameters α, β, a, and b.

The scheme proposed here is to replace the computationally expensive PIC simulation of the plasma with a time-dependent 3-D electric field map of the plasma—a field that can be generated from an analytic expression. In this way, the effect of the unstable plasma can be investigated by tracking proton pencil-beams through the electric field map which typically requires only a few minutes. A wide parameter space was studied by changing the field map of the plasma and rerunning the beam-tracking simulation for each set of plasma parameters.

The rotation of a long plasma column around the axis of the lens was modelled using a simplified charge distribution that preserves the most important features of the actual $m_\theta = 1$ diocotron mode. The charge distribution was chosen to be simple enough that the electric field generated by the plasma could be calculated based on Gauss's law. The radial profile used to model the electron density is shown in fig. 4.14 and is given by

$$\frac{n_e(r)}{n_{e,max}} = \begin{cases} \alpha, & r < a \\ mr + c, & a \leq r \leq b \\ 0, & b < r < 2R_A \end{cases} \tag{4.3}$$

where m, c are functions of α, β, a, b given that $n_e(a) = \alpha \cdot n_{e,max}$ and $n_e(b) = \beta \cdot n_{e,max}$; and r is the radial coordinate measured from the centre of the plasma column. Here, b acts as a cut-off radius—its value is always set such that the electron density drops to zero where the plasma column touches the inner surface of the anode if an offset is introduced between the central axis of the plasma column and the central axis of the lens. Thus, the electron cloud maintains the cylindrical symmetry for any radial offset of the plasma column. Typical values of the parameters that define the electron density above are: $n_{e,max} = 1 \times 10^{14} \, \mathrm{m^{-3}}$, $a = 1 \, \mathrm{cm}$, $\alpha = 1$, $\beta = 0.1$ and a radial offset of the plasma column of $0.5 \, \mathrm{cm}$. Neglecting the conducting surface of the anode and the longitudinal component of the field at the two ends of the plasma, the radial electric field is obtained from

equation (4.3) as

$$E_r = -\frac{en_{\text{e.max}}}{2\epsilon_0} \begin{cases} \alpha r, & r < a \\ \frac{2}{r}\left[\frac{\alpha a^2}{2} + \frac{m}{3}\left(r^3 - a^3\right) + \frac{c}{2}\left(r^2 - a^2\right)\right], & a \leq r \leq b \\ \frac{2}{r}\left[\frac{\alpha a^2}{2} + \frac{m}{3}\left(b^3 - a^3\right) + \frac{c}{2}\left(b^2 - a^2\right)\right], & b < r < 2R_A \end{cases} \tag{4.4}$$

where r is the position vector from the centre of the plasma column.

Equation (4.4) was used to generate a 3-D Cartesian electric field map with zero azimuthal and longitudinal field components everywhere inside the lens. The field map was made time-dependent by introducing a rotation of the axis of symmetry of the plasma around the central axis of the lens with a constant period. In other words, the full field map consisted of a set of values of the radial electric field on a regular 3-D grid inside the volume of the lens replicated for an integer number of time steps. From the particle-in-cell simulations presented above, the plasma column was observed to rotate around the lens axis with a period of the order of $100\,\text{ns}$. Both the total time interval for which the field map is specified and the spatial and temporal resolutions are limited by the computation time of the beam-tracking simulation. The total time interval for which the field map is generated can be increased if the spatial resolution is decreased, keeping the total memory size required by the field map constant. The field map was generated to include at least three and up to ten periods of rotation by modifying the spatial resolution accordingly. Increasing the resolution of the field map increases the time it takes to generate the field map as well as the duration of the beam-tracking simulation. Thus, the time step was fixed at $dt = 0.2\,\text{ns}$ and the spatial grid separation at $0.5\,\text{mm}$. No significant change was observed in tracking the proton beams for smaller time steps or better spatial resolution.

The proton pencil-beams were tracked with BDSIM [141] as the code allows the use of time-dependent field maps and provides a 4-D cubic interpolator to calculate a more realistic electric field at the location of a proton from the values specified on the regular grid of the field map. The interpolator is also applied in calculating the field at points in time between the time steps specified in the field map. The beam-line created in BDSIM has two main elements: a $36\,\text{cm}$-long volume filled with the field map of the plasma inside the $50\,\text{cm}$-long lens and a drift space of $67\,\text{cm}$ downstream of the lens. The protons travel from the entry plane of the lens to the end of the beam-line where their position is recorded to a file. The length of the field map was chosen to be approximately equal to the length of the plasma column extracted from the PIC simulations of the prototype Gabor lens. A separate static magnetic field map was added to account for the marginal focusing effect of the external magnetic field of the lens on top of that of the plasma.

The input proton distribution was loaded into BDSIM by using a file that was generated separately. The file specifies the (x, x', y, y', t, E) coordinates of each particle at the input plane. The x, y coordinates were generated to fill uniformly six circular holes of radius $r = 1\,mm$. The positions of the holes correspond to the aperture plate placed upstream of the Gabor lens during the beam test. The protons have a kinetic energy of $1.4\,\text{MeV}$ and the duration of the beam spill was set to $300\,\text{ns}$ such that it samples several rotations of the electron
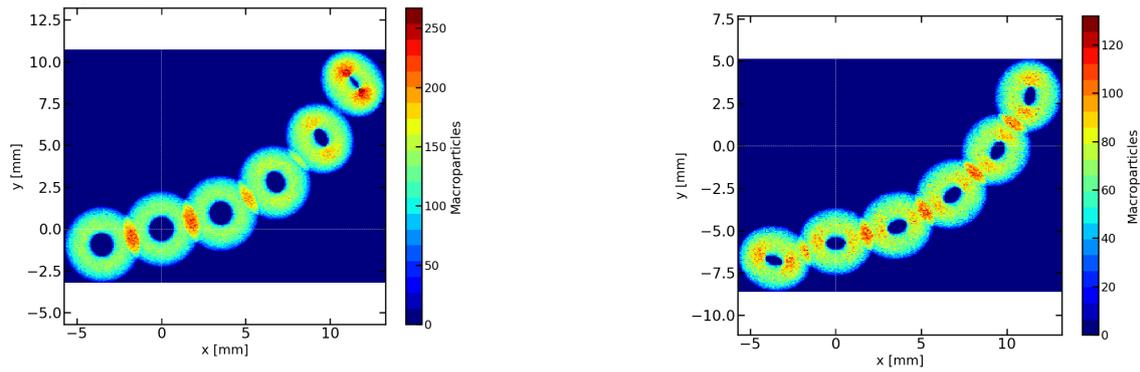
Figure 4.15: Number of macro-particles hitting a screen downstream of the lens for an idealised electron plasma column which rotates around the axis of the lens. The separation between the rings and the eccentricity of each ring increases further away from the axis of the lens. The shape of the rings is also affected by a possible offset between the beam axis and the axis of rotation of the plasma column. The second pencil-beam from the left is aligned (left) with the axis of the lens or the entire pattern of beams is shifted vertically downwards (right) at the entrance of the lens.

cloud. The x' and y' were tuned such that the intensity profile of the six pencil-beams at the screen as obtained with BDSIM matches the images taken during the experiment with the lens off. These profiles were matched in terms of the relative position of the pencil-beams on the phosphor screen, as well as the $x-$ and $y-$diameter of each bright spot.

When the free parameters of the model (maximum electron density $n_{\rm e,max}$, peak width a, radius of rotation, and period of rotation) were tuned accordingly, a qualitative agreement was obtained between the images taken during the experiment (fig. 4.7) for a coil current of 32 A and the results of the simulation shown in fig. 4.15. The eccentricity of the rings increases for the pencil beams that are further away from the beam axis as a result of the different focusing strengths in the x and y directions. The difference in the orientation of the rings in figures 4.7 and 4.15 is caused by the arbitrary definition of the x and y axes on the virtual screen from the simulations. In other words, the intensity pattern observed in simulations is 'viewed' from upstream of the screen, while the camera images from the experiments were recorded from downstream of the screen with the camera pointing towards the plasma lens.

Furthermore, the model predicts the same change in the intensity pattern of the pencil beams as the experiment when the current in the coil is increased from 18 A to 32 A. Similarly to the experiment, fig. 4.16 shows that the rings start to form at a current of 28 A or higher. To obtain the qualitative agreement, the lens was assumed to generate a plasma with a density of 40% of the maximum density, n_B, given by the Brillouin limit. To calculate n_B for each value of the current, the magnitude of the magnetic flux density in the centre of the lens was first evaluated with a simulation of the pancake coils of the lens and, then, converted into the corresponding value of the maximum plasma density.

Figure 4.17 shows that increasing the maximum electron density, $n_{\rm e,max}$, in the central region of the plasma column decreases the separation between the rings and their thickness due to stronger focusing from the plasma. Moreover, the aspect of the spots changes significantly with a small variation of the electron density. Thus,
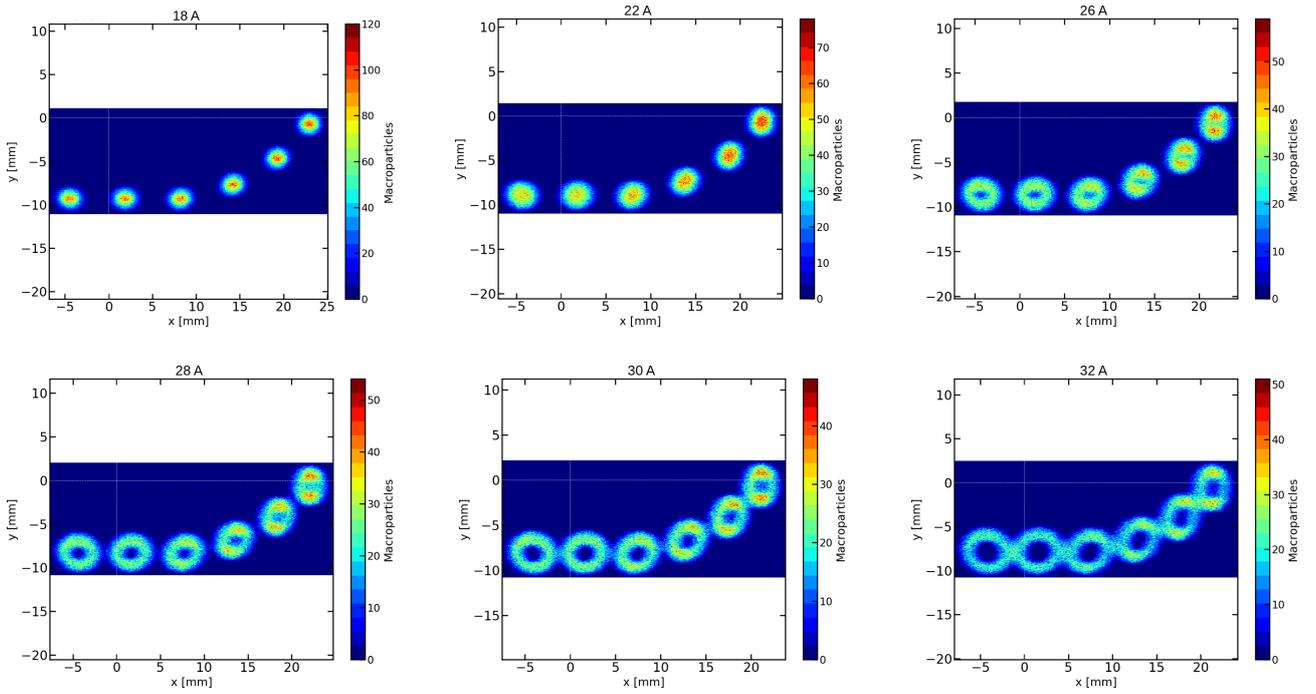
Figure 4.16: Intensity pattern produced by six pencil beams passing through the idealised rotating plasma column on a phosphor screen dowsntream of the lens. The current through the solenoid of the lens was scanned from 18 A to 32 A. A filling factor $\alpha = 40\%$ was assumed for the lens.

matching the pattern from a simulation to an image taken during the experiment provides an approximate indication of the density of the plasma that was confined by the lens during the measurements. However, a similar change in the radius of the rings was obtained by increasing the radius of rotation of the plasma column around the central axis. Consequently, there may exist distinct plasma parameters that produce similar patterns of rings on the screen. Nonetheless, the qualitative agreement between the simulation and the images from the beam test is enough to estimate the order of magnitude of the density of the plasma produced by the lens.

By studying how the rate of rotation of the plasma affects the formation of the rings, fig. 4.18 indicates a lower limit for the rotation period. The corresponding frequency must be slow enough for the rotation of the plasma not to be averaged out during the transit time of the protons. A rotation period which is at least a factor of five higher than the transit time of the protons through the lens is necessary for the formation of complete rings on the screen. As in the experimental observations, the brightness of each ring is seen to vary along its circumference. The rotation period of the plasma that leads to the formation of complete rings in the simulation is also consistent with the numerical prediction from an analytic model of the fundamental diocotron mode. The linearized infinite-length, small-amplitude diocotron frequency is given by [235]

$$f_{\text{dio}}^{\infty} = \frac{eN_L}{4\pi^2\epsilon_0 B R_w^2}, \tag{4.5}$$

which typically provides an estimation smaller by a factor of 2 or 3 compared to the measured value. A diocotron period of the order of $100\,\text{ns}$ is produced in the prototype Gabor lens by a plasma column with a linear electron density, N_L, in the vicinity of $5 \times 10^{11}\,\text{m}^{-1}$ given the configuration of the prototype of the lens. For a plasma

(a) $n_{\mathrm{e,max}} = 1.6$

(b) $n_{\mathrm{e,max}} = 1.8$

(c) $n_{\mathrm{e,max}} = 2.0$

(d) $n_{\mathrm{e,max}} = 4.0$
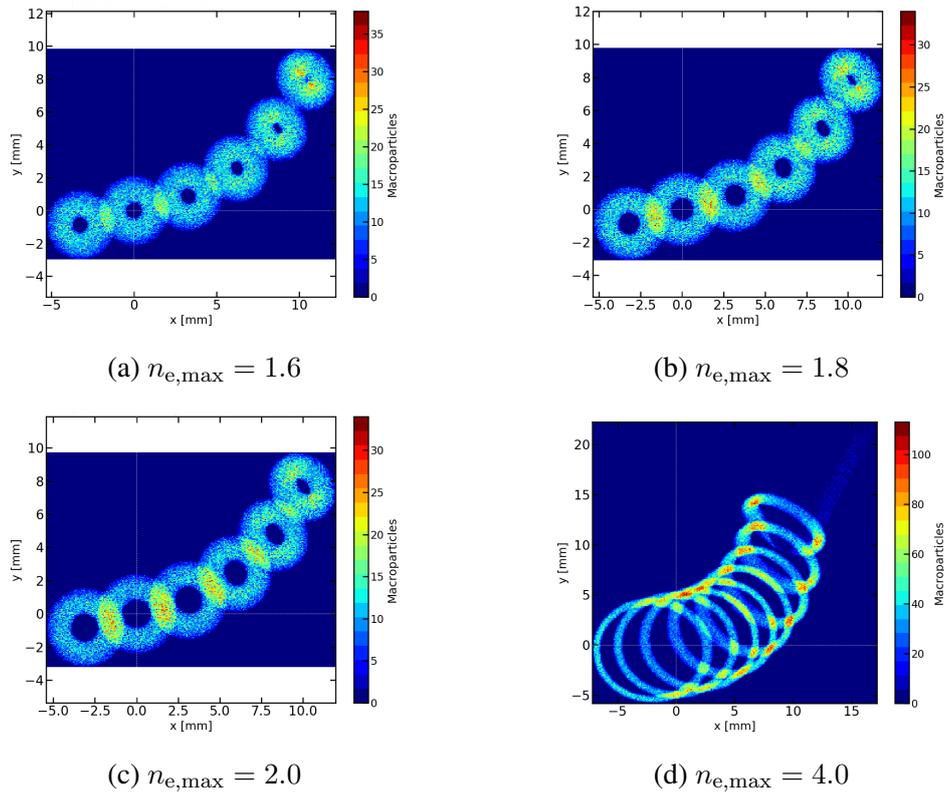
Figure 4.17: The change in the aspect of the rings produced by the proton beamlets onto the phosphor screen with increasing peak plasma density, $n_{\mathrm{e,max}}$ (given in units of $10^{14}\,\mathrm{m}^{-3}$). In case (d), more beamlets were tracked through the model of the lens.

column with a radius as large as $1\,\mathrm{cm}$, the electron density reaches the order of $1 \times 10^{15}\,\mathrm{m}^{-3}$.

If the radius of rotation of the plasma and the filling factor remain constant while changing the current through the pancake coils of the lens, certain trends appear. The linear focusing force is depicted in fig. 4.19 which shows that the centroids of the ring spots move towards the focusing centre for higher electron densities. The six sets of points, each for the corresponding pencil beam, point back to the same focusing centre. This feature does not appear in the experimental data which suggests that the rotation of the plasma column may have been affected when the current in the coil was increased between consecutive measurements. The change in the shape of the rings quantified by the plot on the right of fig.4.19 follows a trend similar to the one recorded in the beam test [54]. However, the result from the measurements is less linear than the prediction from the simplified model in this section which indicates variations of some of the plasma parameters between separate measurements with different coil currents. One possible source of the difference between the predictions of the model and the results of the beam test is the fact that the filling factor of the lens was assumed to be independent of the coil current in the numerical model. In practice, changing the current trough the coil between measurements might have led to variations in the filling factor and, thus, to a non-linear dependence of the peak plasma density on the current in the coil.

In summary, a simple analytic model has been presented above which was used to generate time-dependent electric field maps for rapid scans of various plasma parameters as part of beam-tracking simulations of protons passing through an unstable plasma. It was found that an infinitely long plasma column that is displaced radially

(a) $T = 100\,\text{ns}$ (b) $T = 60\,\text{ns}$ (c) $T = 40\,\text{ns}$
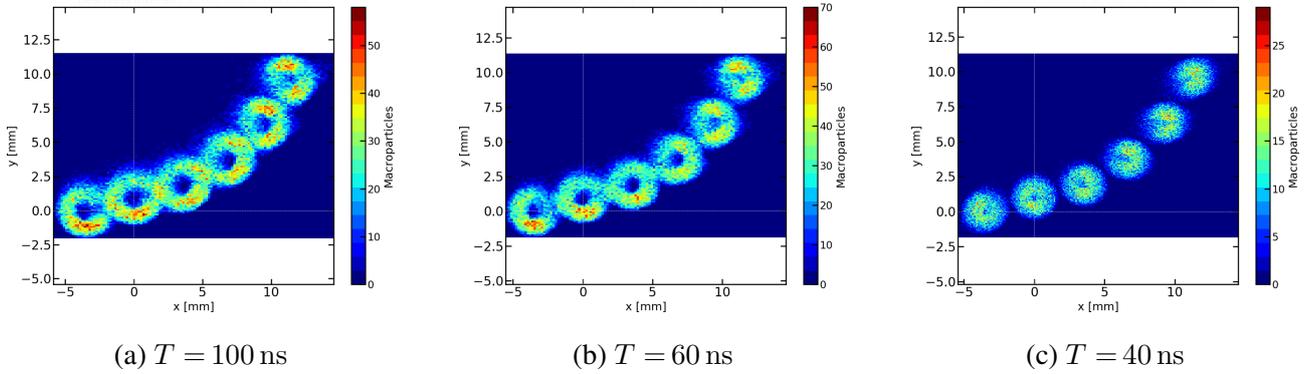
Figure 4.18: The change in the aspect of the rings produced by the proton beamlets onto the phosphor screen for a few values of the period of rotation of the plasma column inside the lens, T. The transit time of the protons through the lens is approximately $20\,\text{ns}$.
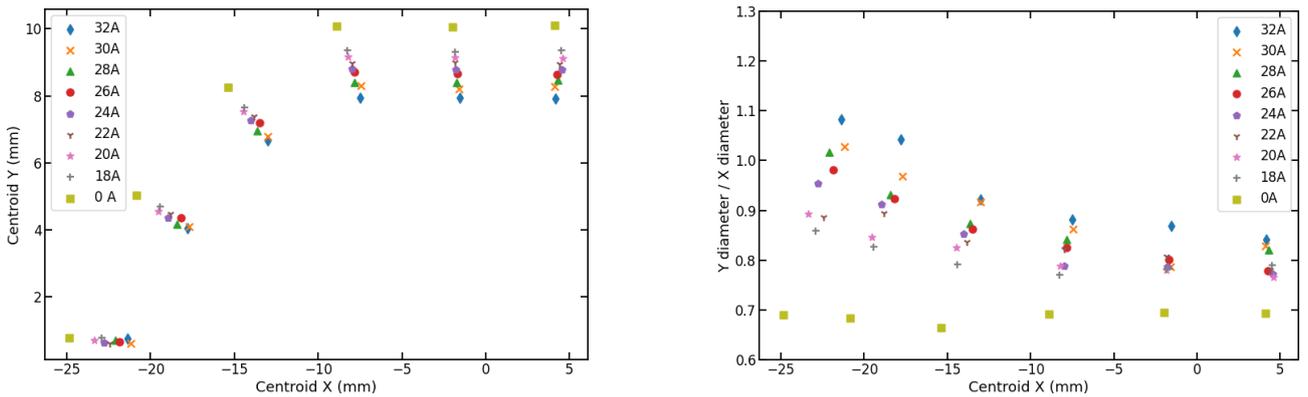
Figure 4.19: **Left:** The position of the centroid of the six rings left by the pencil beams onto the screen. **Right:** The ratio of the y-diameter over the x-diameter of the six rings. The results of the proton-tracking simulations are shown for multiple currents through the coil of the lens.

from the central axis of the lens and rotates around the same axis converts thin proton pencil-beams into rings. The aspect of the rings and the qualitative changes in the aspect for a range of magnitudes of the plasma density, radial off-set, and period of rotation are highly similar to the experimental observations from the beam test of the prototype Gabor lens [54]. Consequently, the $m_\theta = 1$ diocotron mode was identified as the most probable plasma instability that may have caused the anomalous focusing effect observed during the beam test of the lens.

4.3.4 PIC simulation of the plasma $m = 1$ diocotron mode

In the last part of this chapter, a selection of results are presented from PIC simulations that reproduced the $m_\theta = 1$ diocotron rotation of an electron plasma column confined in a Penning-Malmberg trap. The frequency of the diocotron rotation can be measured experimentally relatively easy by picking up the induced-charge signal on azimuthal segments of a cylindrical electrode which are electrically isolated from each other (see, for example, [217] or [236]). Thus, a comparison of the diocotron frequency predicted by a PIC simulation and the value measured experimentally can act as a validation of the PIC code of interest.

In practice, there are several techniques by which electrons are injected and captured inside a Penning-Malmberg trap and form a confined plasma. However, since the time scales involved in the capture processes (e.g. electron collision period, axial bounce period, or stream instability rise time) are much larger than the plasma time scales (e.g. plasma period, cyclotron period, or diocotron period), it is computationally too expensive to simulate the capture of electrons and the subsequent evolution of the plasma towards a stable state inside the lens. The alternative is to initialise the electrons directly in a state as close as achievable to that of a stable confined plasma. A 'hybrid' scheme was devised in this work to initialise the electrons in a PIC simulation such that the plasma remains stable after the start of the simulation. The spatial distribution and the angular velocity of the electrons inside the model of the plasma trap are calculated numerically from the finite temperature theory of the plasma. The central assumption here is that, for small radial displacements of the column from the axis of the trap, the density distribution is not significantly different from the distribution of a plasma column in thermal equilibrium and with no offset from the axis of symmetry. For the interested reader, the complete scheme for initialising the macro-particles is explained in appendix F.2. The scheme allowed better control over the density distribution of the plasma columns modelled in the simulations compared to the case of using simpler methods of initialising the macro-particles in phase-space.

Figure 4.20 shows an example of the evolution of a plasma column that is initially displaced radially from the central axis of the trap in a PIC simulation. The electrons rotate around the longitudinal axis of symmetry of the plasma. Simultaneously, the plasma column rotates around the axis of the trap. The main parameters of the simulation are specified in table 4.2. An electrostatic solver was used and the electrodes of the trap were modelled as perfect electric conductors. The period of the diocotron mode can be extracted from the change in the transverse position of the centre of mass of the plasma with time or from the variation of the electric field at the centre of the trap. Figure 4.21 shows how the two transverse components of the electric field at the centre of the trap vary in time. A best-fit sinusoidal function was applied to extract the diocotron frequency.
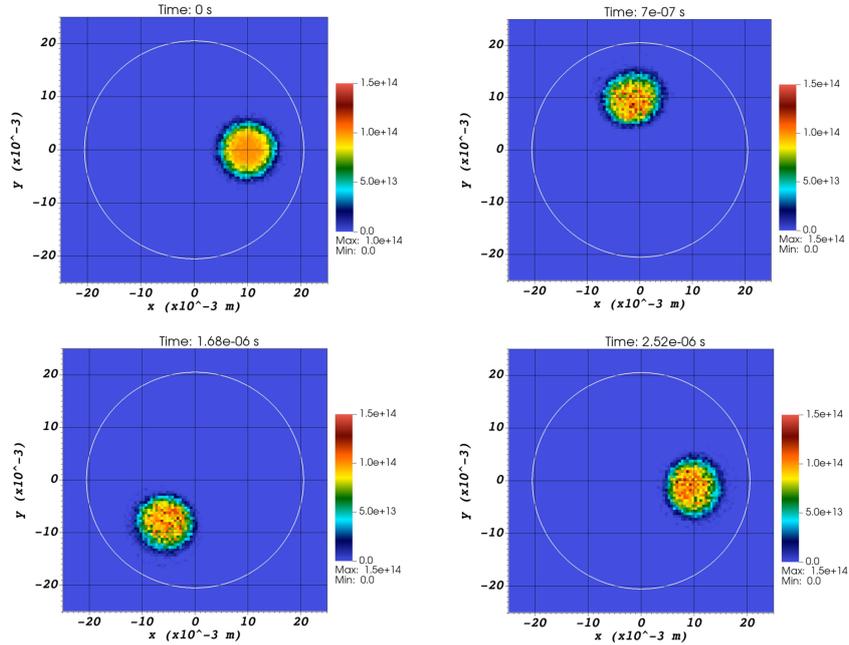
Figure 4.20: Transverse cut through the mid-plane of a plasma lens at four time steps of a PIC simulation. The electron density is indicated by the colour map. The plasma column rotates around the symmetry axis of the trap with an initial displacement of $1\,\mathrm{cm}$. The conducting inner surface of the anode is drawn as a white circle.
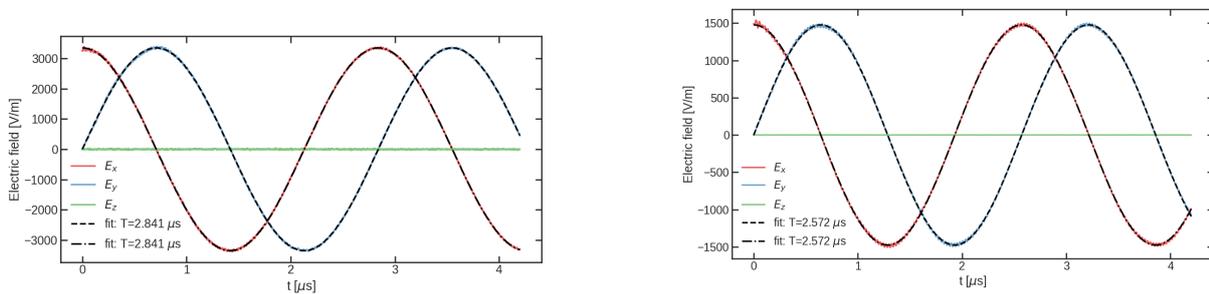
Figure 4.21: The Cartesian components of the electric field extracted at the centre of the plasma trap as a function of time. The plasma column was off-set radially at the beginning of the simulation by $0.5\,\mathrm{cm}$ (left) and $1\,\mathrm{cm}$ (right). A sinusoidal fit is applied to the transverse components of the field to extract the period of rotation of the plasma column around the central axis.

Table 4.2: Main parameters of the simulation of the $m_\theta = 1$ diocotron mode for an electron plasma in a Penning-Malmberg trap.

Parameter	Value
Magnetic field density	$40\,\mathrm{mT}$
Anode potential	$0\,\mathrm{V}$
End electrodes potential	$-150\,\mathrm{kV}$
Electrode inner diameter	$41\,\mathrm{mm}$
Peak electron density	$1 \times 10^{14}\,\mathrm{m}^{-3}$
Plasma temperature	$1\,\mathrm{eV}$
Time step	$\Delta t = 0.14\,\mathrm{ns}$
Grid size	$25\,\mathrm{mm} \times 25\,\mathrm{mm} \times 121\,\mathrm{mm}$
Cell size	$\Delta x = \Delta y = \Delta z = 0.2\,\mathrm{mm}$
Debye length	$0.74\,\mathrm{mm}$
Macro-particles per cell	64

The duration of a PIC simulation that reproduces several diocotron periods may reach the order of a few days on a relatively small number of cores (e.g. 16 CPUs). Therefore, several test cases were run to evaluate the impact of the size of the spatial grid and the time step on the diocotron period extracted from the simulations. Figure 4.22 shows that the temporal resolution has little effect on the extracted frequency for values of the time step that are less than the cyclotron period of the electron and less than the inverse of the plasma frequency. The investigation was not extended to time steps larger than the cyclotron period to avoid introducing nonphysical electron trajectories in the simulation. Currently, VSim does not have a method to integrate corectly the guiding-centre component of the motion of an electron in a magnetised plasma when the cyclotron period is not properly resolved. However, the size of an individual cell in the simulation grid was found to have a greater impact on the diocotron period, as shown in fig. 4.22. A variation of two to three percent in the diocotron period was found when the size of the cell along the longitudinal direction, z, was changed from one Debye length, λ_D, to $0.27\lambda_D$. The diocotron period was found to be less sensitive to changes in the size of the cell along the two transverse directions. Thus, in most of the simulations presented in this work, the cell size was set to the smallest value possible that did not lead to exceedingly long computation times, typically no longer than 48 hours.

Finally, the single example presented above is part of a series of simulations of the diocotron mode reproduced in an electron plasma trap. The diocotron period was extracted from simulations that reproduced typically two periods of rotation of the plasma around the axis of the trap for moderate computational requirements (~ 10 hours per period on 16 CPUs). The most important qualitative features of the $m_\theta = 1$ diocotron mode were observed in the PIC simulations. For example, the diocotron frequency was found to increase with the amplitude of the rotation and the shape of the plasma in the transverse plane was observed to be distorted from circular to elliptical when the plasma is closer to the conducting wall of the trap electrode.
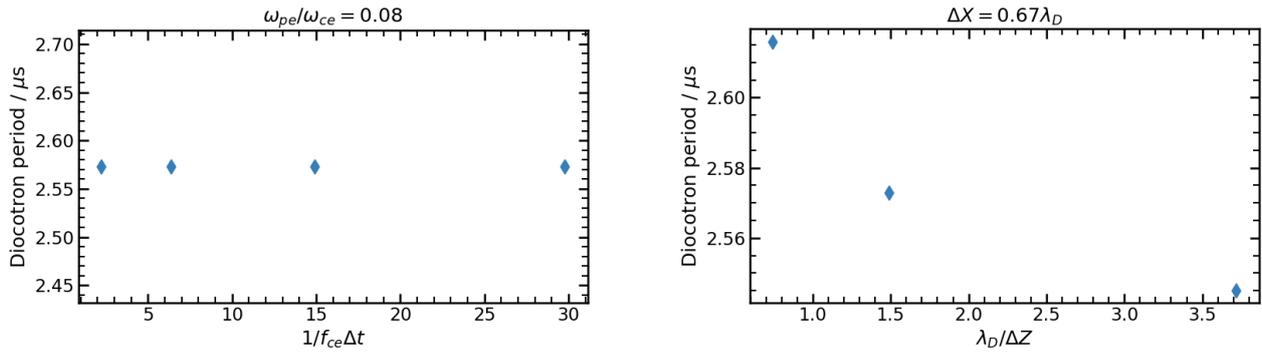
Figure 4.22: The dependence of the diocotron period on the time step (left) and longitudinal cell size (right) in PIC simulations of an electron plasma. The time step is normalised to the cyclotron period ($1/f_{ce}$) and the cell size to the Debye length (λ_D).

4.4 Conclusions

The simulations presented in this chapter were performed to qualitatively understand the unstable plasma regime that could produce deviations of the focusing force produced by a space-charge lens from a linear, aberration-free focusing regime. Another aim of the simulations was to assess whether a commercially available PIC code is suitable for making predictions of the stable and unstable behaviour of the plasma trapped within a space-charge lens. It was demonstrated that with the initial state of the plasma properly defined, several distinct unstable plasma modes can be reproduced with the PIC method.

Concretely, the **diocotron mode with the lowest transverse mode number was identified as a possible source of the anomalous focusing effects** recorded with the Gabor lens prototype built at Imperial College London and tested with proton beams. The starting hypothesis was that the focusing effects observed experimentally are connected to the development of a plasma instability inside the lens. Three most common plasma instabilities were reproduced with the PIC method and their focusing effect on proton pencil-beams was investigated with a particle-tracking code. Of the three unstable plasmas explored, the electron distribution with the largest deviation from the cylindrical symmetry of the trap was found to convert the pencil-beams into rings similar to the ones recorded during the beam test. The rotation of the plasma around the axis of trap in addition to the symmetry breaking of the focusing force is a critical element that leads to the formation of rings. Consequently, a simplified model of a plasma column that rotates around the axis of the lens was constructed. **The model reproduced the main qualitative features of the images of the proton beams that resulted from the beam test.** Furthermore, the agreement between the results of the model and the experimental test was achieved over a restricted range of the density and period of rotation of the electron plasma. Thus, the model provided an approximate evaluation of the plasma density produced and confined by the prototype Gabor lens in the vicinity of $2 \times 10^{14} \, \mathrm{m}^{-3}$ and indicated that the frequency of rotation of the plasma was of the order of $10 \, \mathrm{MHz}$.

Apart from the issue of understanding the capabilities of the Gabor lens prototype, challenges related to the initial distribution of the electrons were encountered during the PIC simulations of both stable plasmas and of the diocotron mode. Transient stages were observed in simulations during which the plasma evolves from the initial

state to a distinct state that remains stable for a longer time. In most of the cases, perturbations in the plasma density caused by this initial transient stage were present for the full duration of the respective simulations. **In order to reduce the transient phase in the evolution of the plasma in a PIC simulation, a 'hybrid' scheme was constructed and tested.** The method relies on the assumption of global thermal equilibrium applied to calculate a realistic 3-D distribution of the plasma density inside the lens. Subsequently, the cold-fluid plasma theory is used to define the radial velocity of the electrons which ensures that the initial plasma does not contract or expand significantly after the start of the simulation. This scheme was shown to work in the case of simulations of the diocotron rotation of an electron plasma column.

Concerning the simulation of the diocotron mode with the lowest mode number in a plasma trap with the Penning-Malmberg configuration, the PIC method was effectively employed to extract the diocotron frequency for a few different plasma parameters. The simulations showed that the diocotron mode can be reproduced with VSim for electron densities of the order of $1 \times 10^{14} \, \text{m}^{-3}$ and plasma lengths of up to $15 \, \text{cm}$ with moderate computational requirements.

It should be noted that most of the results presented in this chapter are qualitative and mostly show the capabilities of the PIC simulation method in the study and development of non-neutral plasma lenses. Some of the simulations given as examples above provide figures of merit that can be compared against measurements obtained with an actual plasma trap that can operate at high electron density. Thus, the development and testing of such a prototype lens is a critical step in the full validation of the simulation methods employed in this work.

In conclusion, the stability of the plasma is a critical requirement for a non-neutral electron plasma lens to act as a reliable focusing device in future proton or ion accelerator beam-lines. The spontaneous formation of the plasma in a Gabor lens was associated with the onset of plasma instabilities that, in turn, lead to time-dependent non-axisymmetric beam-focusing forces. The uniform injection of electrons inside the lens or active control of the plasma rotation after its formation may provide possible solutions to avoid the growth of the most common types of plasma instabilities. For the development of a lens that can produce a highly-uniform, stable plasma, PIC simulations are a valuable technique to undergo preliminary investigations of the focusing capabilities of the lens. Nonetheless, experimental verification is needed to validate the properties of the plasma confined by a lens, once a design has been thoroughly tested in simulations. Part of the results of the studies presented in this chapter have also been published in the journal *Applied Sciences* [54].

As part of the development of a new design for a plasma lens, the issues to be confronted for the production of a more reliable lens were found to be well-known within the scientific community that has been studying low-energy, non-neutral plasmas. However, experiments elsewhere were typically set to investigate plasmas in regimes different from that required by the plasma lenses in LhARA. The following chapter describes a series of experiments with confined electron clouds of comparatively low density. Several preliminary measurements of the parameters of interest for a plasma lens required for LhARA are presented.

Chapter 5

Experimental characterisation of a low-density electron plasma trap

This chapter describes a series of measurements of low-density non-neutral electron plasmas confined by a Penning-Malmberg trap with a similar configuration to that of a space-charge lens. The main goal of the experiments was to explore the electron densities that can be achieved within one of the existing storage traps at Swansea University. In addition to identifying the limitations of the existing trap, the measurements were planned to characterise the properties of the trapped particles such as the typical radial profile, transverse size, lifetime, expansion rate, and the frequency of the transverse rotation of the plasma. Measuring the frequency of the transverse rotation of the electron clouds had been planned as a study of either the magnetron (single-particle regime) or diocotron (coherent plasma regime) motion as such measurements been identified as a means of benchmarking the particle-in-cell software used in previous work.

In the first part of this chapter, an introduction of the motion of electrons inside the trap is given (section 5.1). Then, the experimental setup that was used for the measurements (section 5.2) is described and the methods through which the ensembles of electrons were injected, captured, and ejected from the trap are briefly outlined. The results of the measurements are presented in sections 5.3.1 to 5.3.4 based on the principal characteristics of the confined plasmas which were investigated. Lastly, details are given about simulations of the magnetron motion of the electrons in the trap and the results are compared to data from the measurements (section 5.3.5). The chapter is concluded by a final discussion (section 5.4).

5.1   Electron traps

In preparation for the measurements presented later in this chapter, section 5.1 introduces some relevant theoretical aspects of single-particle dynamics inside an electron trap. The particle trap that was employed for the work presented in this chapter is part of a beam-line for studies with positrons. Given the low number of positrons that are typically stored, the design of the trap is based on the working principles of the Penning-Malmberg trap. This subsection contains an overview of the single-particle dynamics of trapped electrons and positrons in an

ideal Penning trap in order to outline the theoretical basis for some of the experimental observations presented later in the chapter. The main characteristics of the motion of the trapped particles in the ideal Penning trap also extend to the particle dynamics inside the Penning-Malmberg trap which allows easier loading and releasing of the particles. While a more detailed overview of the principles of plasma confinement in a Penning-Malmberg trap is presented in section 3.1, a short discussion on the motion of single electrons inside the potential of an ideal Penning trap is given here.

The ideal Penning trap

The principle of operation of an ideal Penning trap is based on Earnshaw's theorem which states that 3-D static trapping of point charges is not possible by electrostatic forces alone. In order to statically hold a charged particle in a stable equilibrium, a quadrupole potential of the form

$$\phi(x, y, z) = Ax^2 + By^2 + Cz^2 \tag{5.1}$$

is needed with a 3-D minimum ensured if A, B and C are all positive. However, Laplace's equation $\nabla^2 \phi = 0$ requires that $A + B + C = 0$. There are two solutions to achieve stable confinement. One is to use rapidly oscillating electric fields (RF Paul trap) to achieve dynamic trapping. The second one is to apply an axial magnetic field (Penning trap).

In the case of the ideal Penning trap, a ring electrode and two end caps shaped as hyperboloids of revolution generate a pure quadrupole potential. Ions are trapped in the axial direction by a DC potential V_0 applied between the end caps and the ring. The form of the desired potential is

$$\phi(r, z) = V_0 \frac{2z^2 - r^2}{r_0^2 + 2z_0^2}, \tag{5.2}$$

where $2z_0$ is the minimum separation between the end caps and r_0 is the minimum radius of the ring electrode. Equation 5.2 describes a potential with a saddle point at the origin which forms a minimum in the z-direction. Consequently, the radial motion is unstable and a homogeneous axial magnetic field, $\mathbf{B} = B\hat{\mathbf{z}}$, is added to provide radial confinement.

In the superposition of the electric field, $\mathbf{E} = -\nabla \phi$, and the magnetic field, $\mathbf{B}$, a charged particle is subject to the Lorentz force $m\ddot{\mathbf{r}} = e(\mathbf{E} + \mathbf{v} \times \mathbf{B})$. The equation of motion is, then, obtained from

$$m \begin{pmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{z} \end{pmatrix} = \frac{e}{m} \left[\frac{2V_0}{r_0^2 + 2z_0^2} \begin{pmatrix} -x \\ -y \\ 2z \end{pmatrix} + \begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{pmatrix} \times \begin{pmatrix} 0 \\ 0 \\ B \end{pmatrix} \right]. \tag{5.3}$$

Therefore, the single particle equations of motion along each Cartesian direction are

$$\ddot{x} + \omega_c \dot{y} - (\omega_z^2/2)x = 0 \tag{5.4}$$

$$\ddot{y} - \omega_c \dot{x} - (\omega_z^2/2)y = 0 \tag{5.5}$$

$$\ddot{z} + \omega_z^2 z = 0. \tag{5.6}$$

Here, $\omega_c = qB/m$ is the frequency of the cyclotron motion that is associated with particles with charge q in a homogeneous magnetic field. Equation 5.6 shows that the axial motion is decoupled from the transverse motion and corresponds to harmonic oscillation between the end caps with the angular frequency

$$\omega_z^2 = \frac{4qV_0}{m\left(r_0^2 + 2z_0^2\right)}. \tag{5.7}$$

In addition to the axial motion, the $\mathbf{E} \times \mathbf{B}$ field perturbs the cyclotron motion to an orbit described by two separate components: the orbit around a magnetic field line with the modified cyclotron frequency

$$\omega_c' = \omega_c/2 + \sqrt{(\omega_c/2)^2 - \omega_z^2/2} \tag{5.8}$$

and a slow orbit around the trap centre with a frequency ω_m, called magnetron motion, where

$$\omega_m = \omega_c/2 - \sqrt{(\omega_c/2)^2 - \omega_z^2/2} \tag{5.9}$$

The full 3-D path of a charged particle confined in a Penning trap can be seen in fig. 5.1 together with each of the three components of the motion. The centre of the modified cyclotron orbit, referred to as the guiding centre, describes the magnetron motion when projected onto the xy plane. The combined modified cyclotron and magnetron motion traces an epitrochoid in the transverse plane. For typical Penning traps that confine electrons and positrons, the three motional frequencies conform to the hierarchy

$$\omega_c' \gg \omega_z \gg \omega_m. \tag{5.10}$$

For an end cap voltage around $100\,\mathrm{V}$ and trap dimensions close to $r_0 \sim 2\,\mathrm{cm}$ and $z_0 \sim 5\text{–}15\,\mathrm{cm}$ typical values for $f_z = \omega_z/(2\pi)$, $f_c = \omega_c'/(2\pi)$ and $f_m = \omega_m/(2\pi)$ are $\sim 5\text{–}25\,\mathrm{MHz}$, $\sim 1\,\mathrm{GHz}$ and $\sim 5\text{–}500\,\mathrm{kHz}$, respectively.

It is important to note that the average potential and kinetic energy of the magnetron motion are not equal. Thus, the total energy associated with the magnetron motion is negative [237],

$$E = -\frac{1}{4}mr_m^2\omega_z^2 + \frac{1}{2}mv_m^2 = -\frac{1}{2}mr_m^2\omega_m(\omega_c' - \omega_m). \tag{5.11}$$

The consequence is that the magnetron orbit is unstable. Any damping process, such as collisions with the
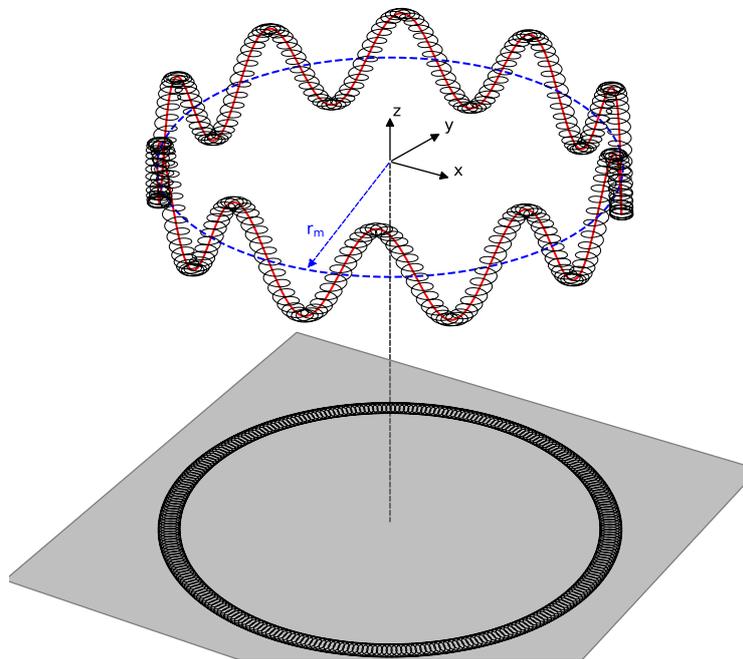
Figure 5.1: The path (black) of a single charged particle confined in an ideal Penning trap. The magnetron component of the motion with radius r_m is shown by the blue dashed line. The "guiding centre" motion is indicated in red and is obtained from the combination of the axial and magnetron motion. On the bottom of the drawing, the path of the particle is projected onto a transverse plane parallel to the x and y axes.

background gas, will cause the trapped particles to go down the potential curve to larger r_m and eventually collide with the walls of the trap. This loss mechanism is referred to as cross-field or radial transport. In practice, a trap has to be held under ultra-high vacuum conditions ($<10^{-9}\,\mathrm{mbar}$) and the trap components must be cleaned appropriately to prevent degassing and contamination of the vacuum. By contrast, the total energy of the axial and modified cyclotron motions are positive and proportional to z_0^2 and r_c^2 respectively, where r_c is the radius of the orbit around the guiding centre. Therefore, allowing the trapped species to interact with an inert buffer gas at room temperature leads to cooling of the modified cyclotron and axial components of the motion.

The Penning-Malmberg trap

The ideal (harmonic) Penning trap does not represent a practical configuration for a plasma lens as the canonical design does not allow for particles to be easily loaded or released from the trap. Therefore, a variation of the Penning trap that is more suitable to ion beam focusing is the cylindrical Penning trap, also known as the Penning-Malmberg trap [188]. In this type of trap, both the ring electrode and end-caps are replaced by a number of co-axial cylindrical electrodes to generate axially confining electric fields (see also fig. 3.1). Thus, the particles to be confined are loaded and unloaded axially, along the magnetic field lines, offering increased versatility that has been proved to be essential in particular for studies of non-neutral plasmas and for experiments with antimatter.

The most important difference between the ideal and the cylindrical Penning trap is the shape of the electrostatic field produced by the electrodes. The potential formed by the cylindrical electrodes is non-trivial to calculate, but an analytical solution [236] can be found by solving Laplace's equation for given boundary conditions. Close to the centre of the trap, and for suitably chosen electrode sizes, the electrostatic potential has a quadratic saddle point similar to that of eq. (5.2). However, there are higher order corrections that lead to deviations from the simple treatment of the magnetron motion above. A first consequence of the non-ideal potential is that the bounce frequency, ω_z, is not independent of the axial kinetic energy of the confined particles, as is the case for the harmonic trap. Furthermore, ω_z also depends on the radial position of the particle and, thus, changes with the amplitude of the magnetron orbit.

5.2 Experimental setup

The experiments described in this chapter were carried out in a section of the positron beam-line at Swansea University labelled as the 'third-stage' which is composed of a Penning-Malmberg trap for both positrons and electrons. The positron beam-line at Swansea University has the distinctive feature of allowing confinement of both positrons and electrons in separate particle traps with diagnostics available for both types of particle. For this reason, a set of initial numerical studies was done to evaluate the suitability of the experimental setup for measurements of the focusing of positron beams by a confined electron plasma. For the interested reader, the main results of the numerical studies are included in appendix E as a reference for future experiments.
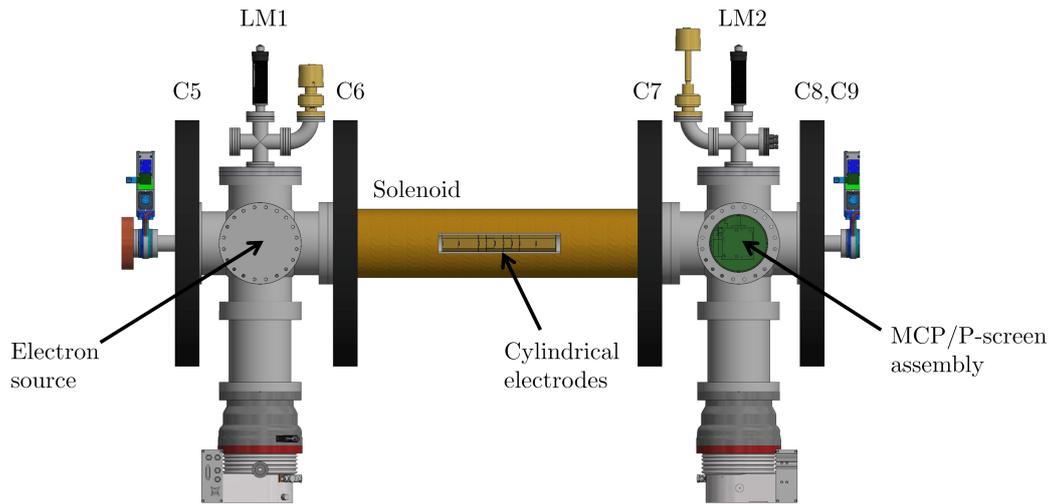
Figure 5.2: The main components of the third stage particle trap of the Swansea positron beam-line. The six-way vacuum cross on the left houses an electron gun mounted on a vertical linear manipulator (LM1). The nine gold-plated aluminium electrodes used for axial trapping of particles are covered by the central solenoid (gold) that traps particles axially. An imaging assembly which contains a microchannel plate (MCP) detector and a phosphor screen is mounted on a second linear manipulator (LM2) inside the vacuum cross on the right. Five pancake coils (C5–C9) guides and focuses the particles when they travel outside of the central solenoid. Each cross is evacuated by a turbomolecular pump.

The measurements described in this chapter were performed solely with confined electrons and no positrons or positron beams.

A description of the complete experimental beam-line dedicated to studies with confined electrons or positrons may be found elsewhere [97, 238, 239]. An overview of the main components of the third-stage particle accumulator is shown in fig. 5.2. The electrons are confined axially in the central region by nine gold-plated aluminium cylindrical electrodes and radially by a water-cooled solenoid that encases the electrodes and the vacuum tube. On either side of the solenoid tube there is a stainless steel six-way vacuum cross, each cross being evacuated by a turbomolecular pump backed out by a scroll pump. Ultra-high-vacuum pressures down to 10^{-9} mbar can be achieved inside the crosses. Upstream of the solenoid, an electron gun is mounted on a vertical linear manipulator inside the corresponding vacuum cross. Downstream of the trapping region, the second cross houses a microchannel plate (MCP) and phosphor screen assembly also installed on an independent linear manipulator. Five large pancake coils are distributed along the beam-line to provide steering and focusing of the particles.

The measurements with confined electrons that are described in this chapter were predominantly obtained with the same sequence of the following three steps: capture of electrons inside the trapping volume, study or storage of the confined electrons, and ejection of the particles towards the diagnostic section of the beam-line. Each of these three stages of the main measurement procedure rely on a particular section of the beam-line that is shown in fig. 5.2. The electrons are supplied as a beam by a thermionic source placed upstream of the trapping section. After the storage and study stage, the electrons are ejected towards a detector assembly situated downstream of the storage trap. The electron source, storage trap, and detector assembly are described in more detail below.
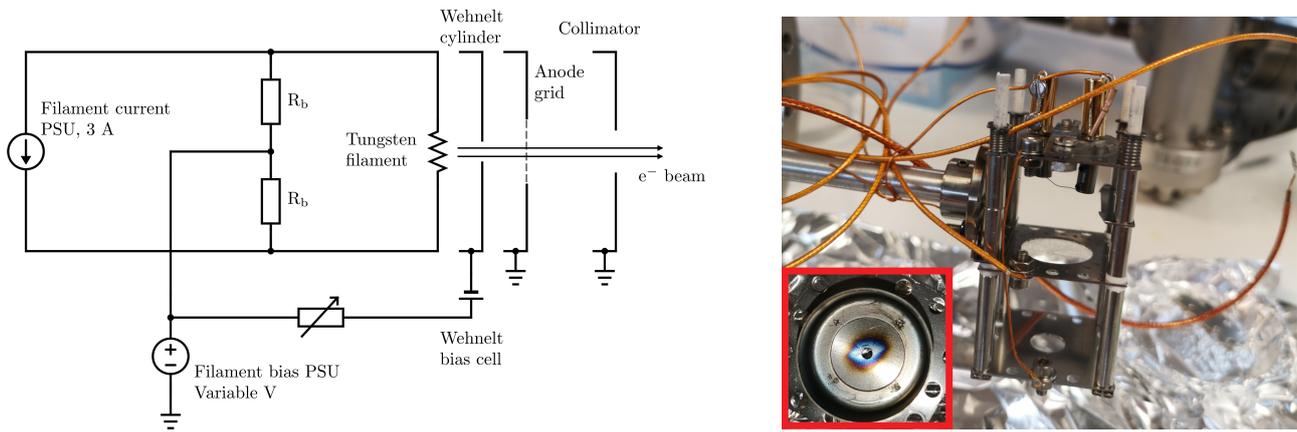
Figure 5.3: Effective circuit diagram of the electron gun used to inject particles into the storage trap (left). Image of the gun (right) dismounted from the beam-line with the Wehnelt cylinder removed and showed in the inset picture on the right. The tungsten filament was typically heated with 3 A and biased at 30 V. The gun produced an electron beam with an emission current of up to about 10 µA. With the Wehnelt cylinder removed, the typical emission current could be increased to several tens of µA with an upper limit of about 300 µA.

5.2.1 Electron source

The electron source is based on the typical structure of a thermionic electron gun and it is schematically shown in fig. 5.3. The electrons are emitted from a thoriated tungsten wire which is kept curved between two electric connectors. The emission current can be adjusted by changing the current from the power supply unit (PSU) connected to the filament. In order to avoid damaging the filament, the current supplied was kept below a maximum value of 3.3 A. The electrons are extracted from the surface of the wire by a fine metallic grid biased to a positive voltage and situated downstream of the filament. The bias voltage of the grid also determines the parallel energy of the electrons as they exit the source.

In order to collimate the electrons and, thus, obtain a more parallel beam, two circular apertures are included in the source assembly. Measurements of the beam current obtained from the source showed that removing the narrower collimator (the Wehnelt cylinder in fig. 5.3) led to a significant increase in the electron emission current up to an order of magnitude. The higher beam current obtained in the absence of the narrow collimator led to a larger number of electrons that were typically trapped by rapidly capturing part of the beam from the source. An example of the number of electrons trapped with the modified source and the beam-capture method is shown in fig. 5.4. To increase the emission current, the filament was typically left on overnight prior to any measurements. The typical emission currents measured on the filament were of several tens of µA with a maximum achievable limit of about 300 µA with the Wehnelt cylinder removed and a current of 3.3 A through the filament. The corresponding electron-beam currents picked-up by the P-screen downstream of the storage trap with all the electrodes grounded was typically between 1 µA and 10 µA with the Wehnelt cylinder and up to 150 µA without the cylinder when the solenoid and the guiding coils were energised in both cases. The beam current measured at the location of the P-screen assembly was highly sensitive to the vertical position of the source and to the alignment of the guiding coils. For currents through the filament between 3.1 A and 3.3 A and,
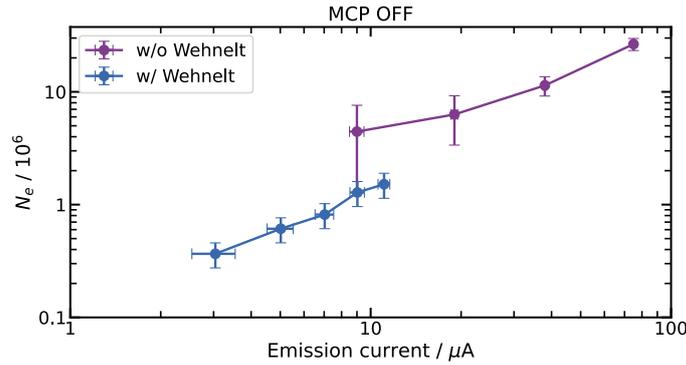
Figure 5.4: The number of electrons trapped with the beam-capture technique over the range of typically achievable filament emission currents for the source with and without the Wehnelt cylinder. The electrons were trapped with the outermost electrodes as end gates.

hence, emission currents above $\sim 100\,\mu$A, two electron columns were observed in the trap. The two electron populations were probably produced by two separate emission spots on the filament. Furthermore, at similarly high emission currents, the images of the phosphor screen show a comet-like tail following the rotation of the main electron cloud in the transverse plane. The tail was formed by electrons spiralling radially outwards from the core of the particle ensemble.

The electron source is installed on a linear manipulator that determines the position of the source assembly along a vertical axis. Thus, the electrons can be injected off axis with respect to the symmetry axis of the storage trap or the source can be better positioned to improve the alignment with the guiding or the trapping magnetic fields.

Electron beam energy

As described above, the energy of the electron beam from the source is dictated by the positive voltage difference between a fine-mesh grid and the filament. A voltage source provided an accelerating voltage of up to $30\,$V between the hot wire and the grid.

The energy profile of the electron beam produced by the source was measured in order to asses the spread of the parallel energy of the electrons that reach the storage trap. Here, *parallel energy* denotes the kinetic energy that corresponds only to the axial component of the velocity of a single particle. During a measurement of the beam energy, all but one of the electrodes of the storage trap were set to ground voltage. The remaining electrode was set to a negative retarding potential that blocked all the electrons with an energy lower than the equivalent decelerating voltage supplied by the electrode. Measuring the current of the beam that passes the decelerating electrode by scanning the magnitude of the voltage applied to the electrode results in a cumulative distribution of the number of electrons as a function of the blocking voltage as shown in fig. 5.5. The energy profile of the electron beam is then calculated from the cumulative energy distribution. It must be noted that this technique only allows the measurement of the axial component of the beam energy. Furthermore, fig. 5.5 shows an enhancement of the transmitted beam current when the decelerating voltage is turned on. This may

result from the negatively biased electrode acting as a beam-focusing lens and, thus, increasing the transmitted number of electrons.

Figure 5.5 shows the energy distribution of the electrons produced by the source for two different extraction voltages. Four separate measurements are averaged for each individual energy bin. Furthermore, the scan over the values of the decelerating voltage was done by shuffling the measurements for each fixed voltage instead of a monotonously increasing scan to avoid the influence of the reflected electrons on the incoming beam from the source. Both measurements show the electron number to peak near the maximum beam energy, which is determined by the acceleration voltage applied to the grid, and a sharp drop-off. The electron beam typically has a wide energy spread of about $20\,\text{eV}$ with lower particle numbers towards the low-energy tail of the distribution. Possible sources of beam-energy broadening are the magnetic mirroring effect and the misalignment of the guiding magnetic fields with the beam-propagation axis.
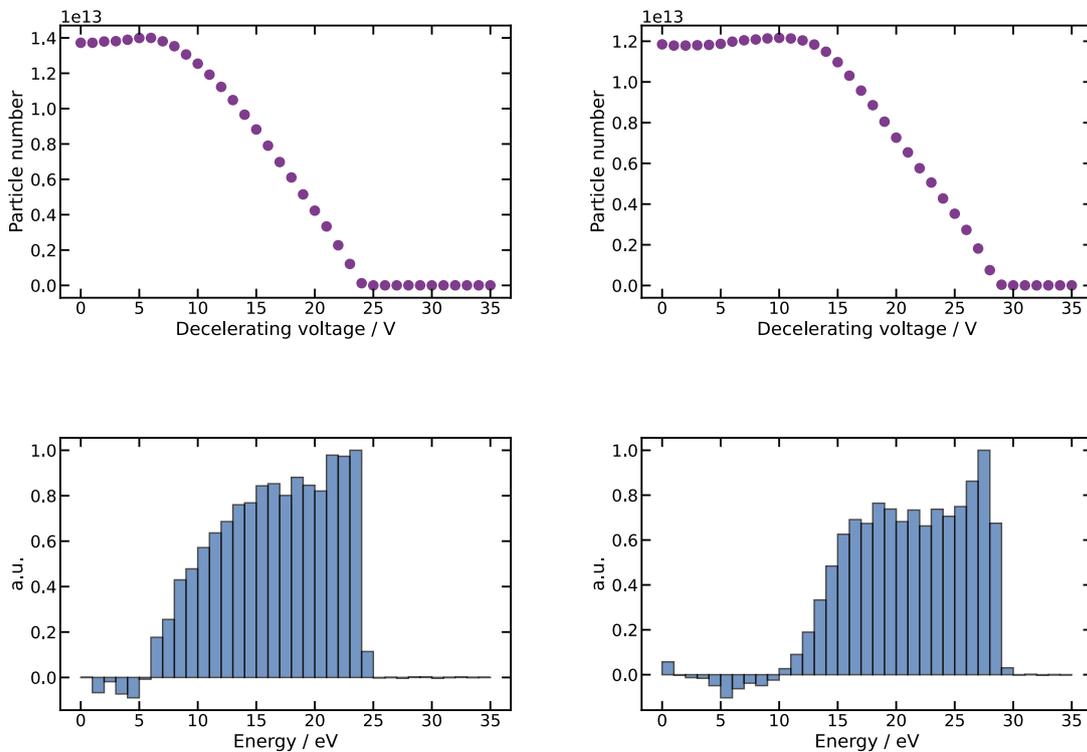
Figure 5.5: **Top:** The number of electrons measured downstream of the storage trap as a function of the decelerating voltage applied to one of the cylindrical electrodes of the trap with respect to the potential of the source. Each data point is the average of four separate measurements. The extraction voltage at the source is $25\,\text{V}$ and $30\,\text{V}$ for the left and right columns, respectively. **Bottom:** The axial-energy profile of the electrons emitted by the source for the two different extraction voltages.

5.2.2 Storage trap

The longitudinal and radial confinement of electrons is achieved with the configuration of confining fields typical to the Penning-Malmberg trap. The fields are generated by a series of cylindrical electrodes surrounded by a solenoid magnet. The electrodes are made from gold plated aluminium, have an inner diameter of $41\,\text{mm}$, can
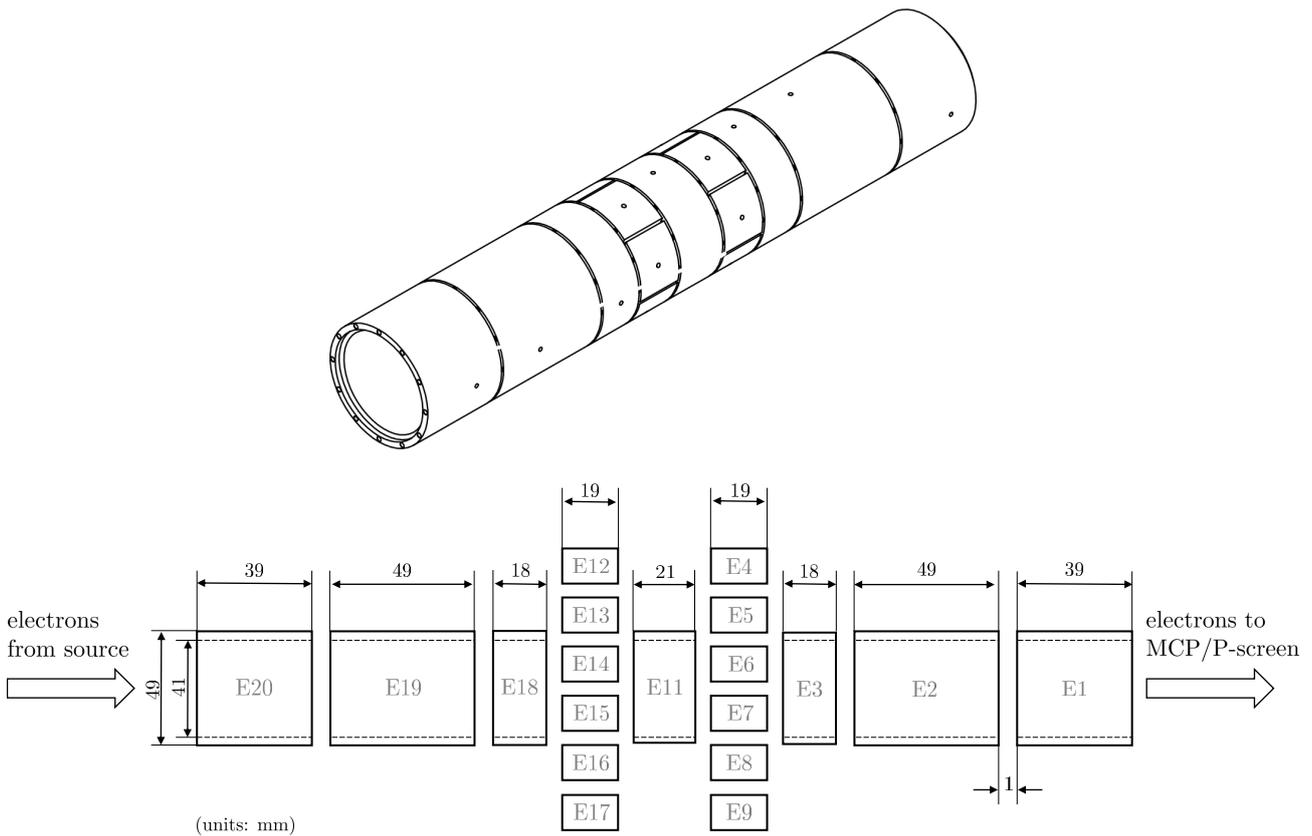
Figure 5.6: **Top**: AutoDesk Inventor™ diagram of the electrodes of the third-stage positron/electron trap [239]. The electrodes are separated by insulating sapphire spheres and each has a hole for the electrical connection. **Bottom**: Schematic of the electrode structure in the storage trap. The electrodes are numbered according to the electrical connections to the external breakout box. Electrodes 4 to 9 and 12 to 17 represent azimuthal segments of two complete ring electrodes designed for application of rotating electric fields.
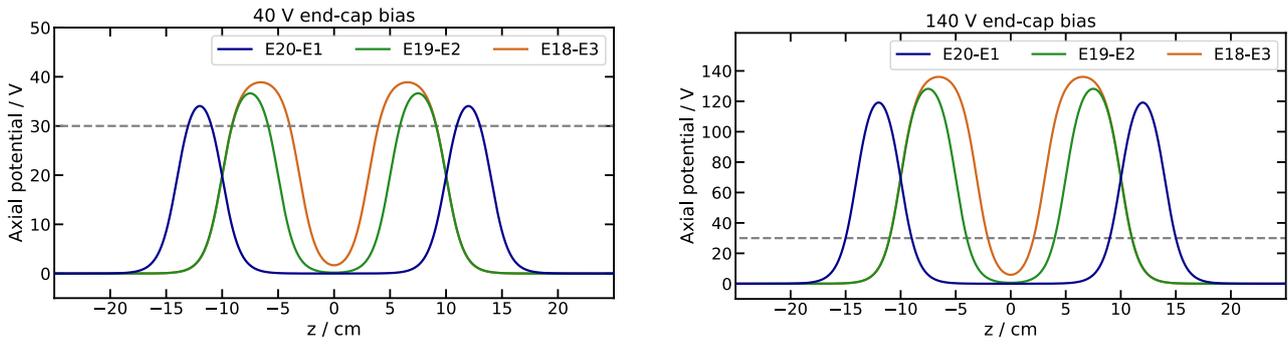
Figure 5.7: The axial electrostatic potential created by applying a bias of $40\,\mathrm{V}$ (left) and $140\,\mathrm{V}$ (right) to the end-cap electrodes labelled E## and E# as indicated in the figures. The dashed line marks the maximum longitudinal energy of the electrons injected into the trap. The electrostatic potential was calculated by solving Laplace's equation for an azimuthally symmetric geometry with the method described in [236].

be individually biased to a voltage up to $140\,\mathrm{V}$, but vary in length as shown in fig. 5.6. Any two electrodes can be used as end caps of the trap and the distance between the two electrodes determines the length of the resulting electron column. The different lengths of the electrodes allow the longitudinal extent of the volume where the electrons can be confined to be varied over several values by using different combinations of electrodes as the two end caps of the trap. Figure 5.7 shows how length of the electrostatic potential well that traps the electrons longitudinally can be varied by choosing different pairs of electrodes to act as end caps. Furthermore, the length of the trapped electron column is also influenced by the shape of the axial potential well and, thus, by the voltage applied to the end-cap electrodes. It must be noted that the axial electrostatic potential well is not quadratic, especially for the longer configurations of the trap. Thus, the motion of the electrons inside such potentials is expected to show deviations from that of an ideal Penning trap.

Two of the electrodes have six azimuthal segments for the application of a rotating electric field as the basis of the so-called 'rotating wall' (RW) technique. By the application of RW the electron columns can be compressed radially [212] through injection of angular momentum into the plasma at a rate that balances the outward expansion [240]. To provide electrical isolation, the electrodes are separated by $2\,\mathrm{mm}$-diameter sapphire balls mounted into indents that are machined into the ends of each electrode creating a gap of $1\,\mathrm{mm}$ between adjacent electrodes. A wire connects each electrode to a breakout box outside of the trap solenoid and is attached to each electrode via titanium screws to limit the interaction with the surrounding magnetic field.

5.2.3   Electron imaging assembly

The main figure of merit for an electron trap working as a plasma lens is the total number of trapped electrons or the corresponding particle density. In a single sequence, the electrons were released from the trap by rapidly lowering the bias voltage of the downstream end-cap electrode. Then, the number of electrons ejected from the storage trap was measured by recording the charge of the particles that strike a conductive surface positioned downstream of the trap. Further information on the position of the electrons and their transverse distribution was

obtained using a scintillating material coupled with a photo-sensitive detector to pick up the photons produced by the electrons that hit the scintillating material. A special assembly was constructed to allow both types of measurement described above. The assembly, illustrated in fig. 5.8, is mounted on a vertical linear manipulator inside the vacuum cross downstream of the electron storage trap. The principal components of the assembly are the micro-channel plate (MCP) and the phosphor screen (P-screen).

The role of the MCP is to amplify the number of electrons that are ejected from the storage trap by several orders of magnitude before they hit the P-screen as the typical total number of trapped electrons is too low to produce a large enough photon shower that can be imaged by a camera. An MCP is a highly resistive plate (typically leaded glass) with millions of channels which pass from the entry face of the plate to the exit face. The channels have a typical diameter of the order of $10\,\mu m$ and are separated by a typical spacing of $12\,\mu m$. An accelerating electric field is created inside each channel by two layers of conducting material deposited on the two transverse faces of the MCP that can be individually biased to high voltages. The channels of the MCP are typically angled with respect to the normal axis to maximise the probability for en electron passing through the plate to hit one of the walls and release other electrons. Thus, the MCP acts as a million continuous dynodes arranged in parallel to multiply the incident electron signal. The large number of collisions within each channel results in a very high gain.

The MCP and the P-screen used in the measurements described in this chapter were supplied by *Lambert Instruments*. The MCP consists of two multi-channel plates with an active area $40\,mm$ in diameter arranged such that the channels of the two plates form a chevron pattern to increase the number of collisions of the electrons with the walls of the channels. The structure and arrangement of the channels, as well as the high isolation between neighbouring channels, help to preserve most of the transverse spatial information of the incident electron distribution. A phosphor screen composed of P43 ($GdO_2S{:}Tb$) is located behind the MCP plates. This material emits light in the green part of the visible spectrum when struck by a charged particle through excitation and relaxation of electron-hole pairs. The phosphor screen was also used as an anode to measure the amplified electron current. A schematic diagram of the imaging assembly and its functions is shown in fig. 5.9.

A mirror mounted at $45°$ behind the screen reflects the scintillation light through a window out of the vacuum system, where it is captured by a *Hamamatsu* Orca-R2 CCD camera, then processed and recorded with Labview™. The camera is placed in a light-tight box to block the ambient light. In front of the MCP, there is an additional fine mesh grid which can be biased to negative voltage in order to protect the MCP/P-screen assembly from a large number or large density of electrons to be accidentally dumped onto it and, thus, to damage the active area of the phosphor screen.

For the work with electrons described below, the back of the MCP was typically biased between $1.5\,kV$ and $1.8\,kV$, while the front was maintained at $0.1\,kV$. The bias on the back plate was adjusted in accordance to the number of electrons dumped onto the MCP in order to avoid saturation of the camera images and simultaneously obtain a high enough gain for the radial distribution of the electrons to be visible. A bias of $2.0\,kV$ or $2.1\,kV$ was maintained on the phosphor screen.
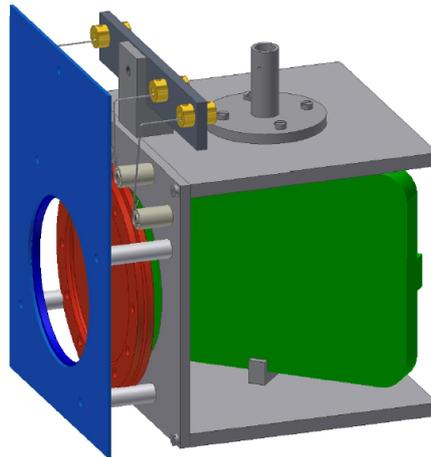
Figure 5.8: AutoDesk Inventor™ rendering of the microchannel plate (MCP) and phosphor screen imaging detector in false colour: protective grid mount (blue), MCP and P-screen assembly (red), mirror (green), HV contacts (gold) and assembly rig (grey). Figure rendered by T. Mortensen, taken from [241].

Experiment control

The studies reported in this chapter mostly employed a capture-store-release cycle, also referred to as a 'sequence', for each individual measurement of certain parameters of the electrons stored in the plasma trap. A more detailed description of a typical experimental sequence employed in the experiments is included for the interested reader in appendix C. A summary of the main aspects of the experimental sequence is given in the following.

Firstly, electrons were loaded inside the volume of the particle trap from the beam emitted by the source via two distinct techniques (appendix C.1). They mainly differ in the extent of the time interval during which electrons were allowed to flow into the confinement region and, as a result, also in the total number of particles and the radial size of the electron clouds that were subsequently trapped. The 'dynamic' trapping technique consisted in a rapid change of the confining potential of the trap such that the part of the electron beam between the two end electrodes of the trap becomes trapped. Thin beam-like electron columns with a comparatively low number of particles were typically produced via this dynamic beam capture. Alternatively, the 'static' particle trapping procedure was based on a slower continuous accumulation of electrons over time intervals of tens or hundreds of milliseconds by partially lowering the bias potential on the upstream end electrode of the trap. Larger number of electrons could be accumulated via the static technique and the resulting electron columns typically had larger diameter due to radial expansion of the space-charge.

Secondly, most of the work presented in this chapter was greatly facilitated by the flexible experimental system that monitors and controls the Swansea positron beam-line. The majority of the experiments performed with the electron storage trap were primarily set up in Labview™ 'virtual instruments' and the corresponding
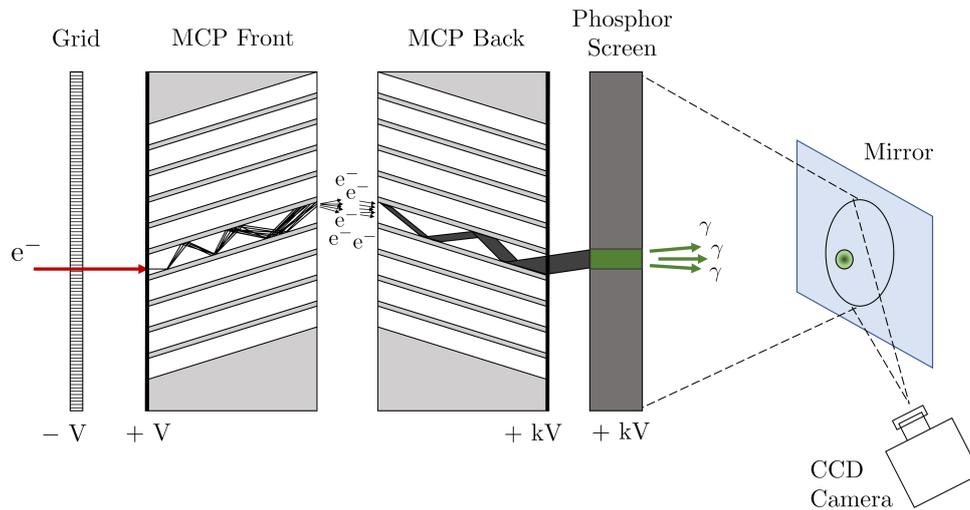
Figure 5.9: Schematic of the detector assembly. An electron that is incident on the biased front surface of the MCP leads to a cascade of electrons through collisions with the walls of a channel. The electron cascade is converted into a photon signal by a phosphor screen. The typical magnitudes of the bias voltages follow P-screen>MCP Back≫MCP Front. A fine mesh grid can be biased to shield the detector from incident electrons. The MCP channel radius, length, and angles are greatly exaggerated for clarity.

graphical user interfaces. More details of the control system can be found in [239] and [242].

A typical experimental run to study the confinement of an electron cloud had three main stages: capture of the electrons from the source into the confinement section of the beam-line, storage of the electron within the trapping electrodes for a specific time interval, and the subsequent ejection of the electrons towards the imaging detector. All these stages require a configuration of the voltages applied to the trap electrodes during each stage and of the time duration of each stage. In many cases, the electrode voltages were also varied within the storage stage of an experimental run. Setting up the automated succession of changes within the particle trap, the timing steps, and triggering the data acquisition at the appropriate time were possible with the use of the so-called 'Sequencer' (appendix C.2). The Sequencer and control system offered a high flexibility to conduct complex and repetitive experimental runs.

Several data acquisition devices were integrated with the control system and allowed the automatic recording of raw data during each individual experimental sequence as set up by the user. Two data acquisition devices (appendix C.3) were predominantly used in this work:

- digitiser—the digitiser was used predominantly to record the voltage response of the charge collector which was subsequently converted into charge to obtain the total number of electrons in the trap;

- CCD camera—the camera recorded the scintillation light from the MCP/P-screen assembly and the resulting images were processed to extract the transverse spatial distribution of the collected electrons.

The voltage traces from the digitiser and the images from the camera were analysed with the use of Python scripts. In addition, the charge amplification factor of the MCP and the capacitance of the charge collectors were separately measured during the experiments (appendix C.4).
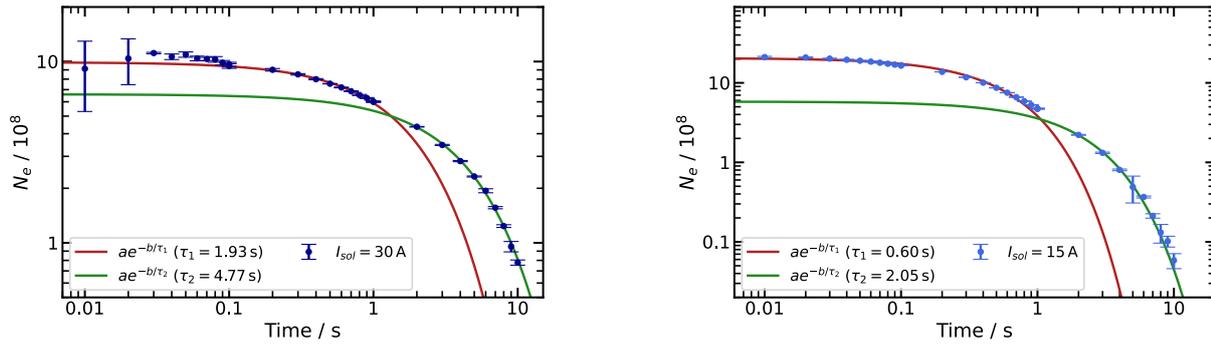
Figure 5.10: Total number of electrons, N_e, released from the storage trap as a function of the duration of time over which they were confined for the trapping solenoid being driven by 30 A (left) and 15 A (right). A stronger current, I_{sol}, through the solenoid of the trap leads to a longer lifetime. The loss of electrons is fitted by a faster exponential function followed by a slower one. The electrons were accumulated via the static technique prior to being stored for a certain amount of time.

5.3   Results

The main results of the experiments with trapped electrons are presented in this section. Some of the quantities were directly measured, while others were derived from the raw data. The magnitude of most of the results, such as the maximum achievable electron density, the transverse size or the lifetime of the electron clouds determine an order of magnitude or a range for the parameters of interest that describe the electron columns that can be trapped and stored with the available apparatus at Swansea. In addition, the frequency of the transverse rotation of the electrons was measured to serve as a means of verifying the simulation model of the trap and of the dynamics of the confined particles. The results presented below are organised by the parameter of interest extracted from the measurements.

5.3.1   Electron cloud lifetime

For an electron plasma to work as a focusing lens for ion beams, the lifetime of the electrons must be significantly larger than the transit time of the ions through the plasma. To avoid refilling a plasma lens after the passage of each ion bunch, it is desirable to maintain the plasma at a constant density for time intervals of the order of seconds. Here, the lifetime of the electrons was measured for the low-density electron columns confined by the particle trap used in these experiments.

In order to measure the lifetime of a large number of electrons, they were trapped via the static technique and using the outmost electrodes as end-caps (E20 and E1). The electrons were accumulated for $200\,\mathrm{ms}$, stored for a certain length of time from milliseconds to seconds, and, then, released towards the MCP/P-screen assembly where the charge was collected and recorded with no bias on the MCP plates. For each individual storage time, the sequence was repeated ten times to obtain the average number of electrons that remained confined. The results are presented in fig. 5.10. The total number of electrons remains nearly constant when the particles are

kept in the trap for less than approximately $100\,\mathrm{ms}$. For longer intervals, the electrons start to be lost gradually and reach half of their initial total number after approximately $1\,\mathrm{s}$ of confinement.

The decay of the total number of electrons with time can be fitted using two exponentially decaying functions, with a shorter and longer lifetime, respectively, as illustrated in fig. 5.10. Furthermore, the two lifetimes were found to increase for a stronger axial magnetic field in the trap. Hence, the decay of the number of electrons measured here suggests that there are two main radial transport mechanisms that cause electrons to be lost on the inner surface of the electrodes. The two mechanisms have distinct time scales and/or are affected differently by the electron density. It must be noted that these lifetime measurements were performed in the absence of any active method for reducing the radial loss of electrons, such as the rotating-wall technique [240].

5.3.2   Electron cloud expansion rate

In the absence of active methods to reduce the the radial loss of particles, the confined plasma expands radially and electrons are lost on the internal walls of the electrodes. The expansion rate of the plasma is usually accompanied by a decrease in the central density of the space-charge column. For an electron plasma to serve as a reliable ion-focusing lens, the central region of the plasma should maintain a quasi-constant space-charge density for the time desired to focus multiple ion bunches. Once the central density drops down, the lens must be refilled before further ion bunches can be focused.

To assess the time it takes for the central electron density to drop to a fraction of the initial magnitude, a capture-store-release sequence was run with variable duration of the interval during which particles were stored. Figure 5.11 shows how the peak light intensity from the camera images drops as a function of time. As the peak electron density in the trap is proportional to the peak intensity of light produced when the electrons hit the phosphor screen, the curves in fig. 5.11 also show the decrease in time of the peak space-charge density of the trapped particles. For the fastest expanding electron ensembles, it took about $50\,\mathrm{ms}$ for the peak density to drop to half of its initial value. Furthermore, the radial size of the electron columns increases at a rate which slows down in time. After about $100\,\mathrm{ms}$. the transverse size of the electron columns becomes comparable to the diameter of the trapping electrodes, indicating that particles start to reach the conducting walls and be lost.

Furthermore, the rate of radial expansion is seen to be dependent on the length of the electron column, L_p. The dependence of the radial expansion rate on the L_p/B ratio is well-known from other plasma-confinement apparatuses [197]. The slower rates of radial expansion were measured for the shorter configurations of the electron trap. A particularly slow rate of decrease of the peak density was measured for $6\,\mathrm{cm}$-long trap. As a conclusion, a particle trap capable to store an electron plasma more than a few tens of centimetres long with a quasi-steady central density for more than $100\,\mathrm{ms}$ would require larger magnetic fields and/or an active mechanism for the reduction of radial transport of the particles.
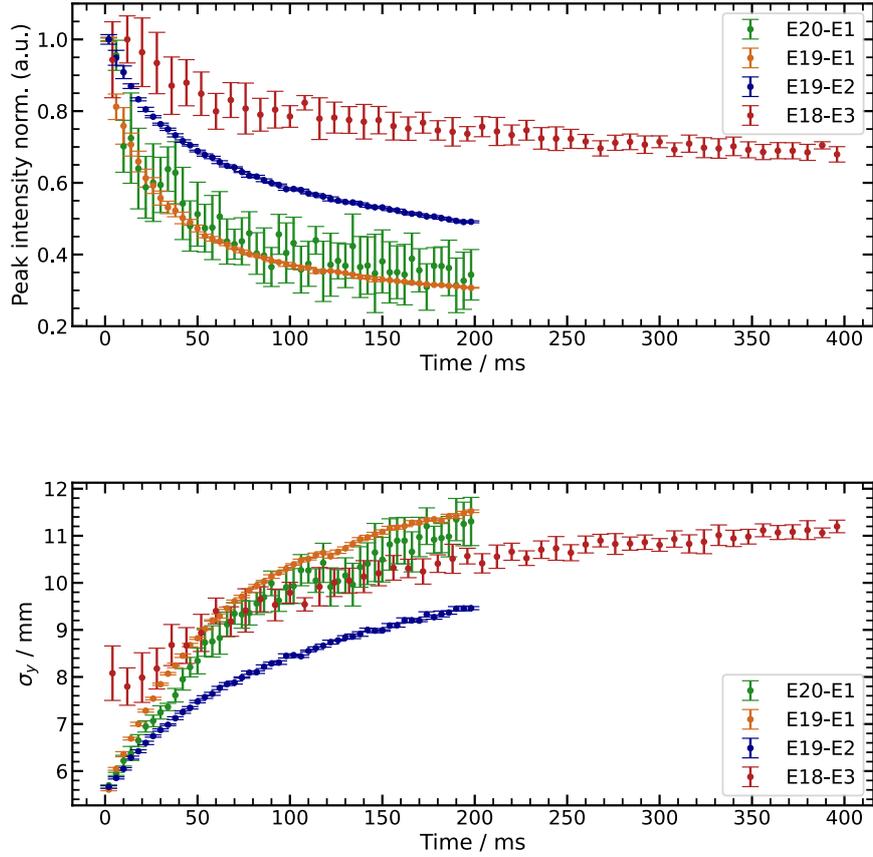
Figure 5.11: **Top:** Normalised light peak intensity measured on the phosphor screen for an electron cloud released from the trap at increasing time intervals. **Bottom:** The radial size of the electron cloud in the trap as a function of the time for which the electrons were confined. The labels indicate which trap electrodes were used as end gates and the approximate lengths of the electron columns are: 20 cm (E20-E1),15 cm (E19-E1),10 cm (E19-E2),6 cm (E18-E3).

Table 5.1: The experimental sequence and electrode voltages used in the study of the maximum space-charge that could be trapped via the static trapping technique (slow electron accumulation). $V_{\min}$ was varied in the range -140 V to -20 V and t_{acc} from 0 to 500 ms.

Phase	Duration	Voltages [V]		
		E20	E19–E2	E1
Reset	100 ms	-30	-30	-140
Open potential well	10 ms	-5	$V_{\min}$	-140
Accumulate	0 ms $< t_{\mathrm{acc}} <$ 500 ms	-5	$V_{\min}$	-140
Release	10 µs	-140	-5	-5
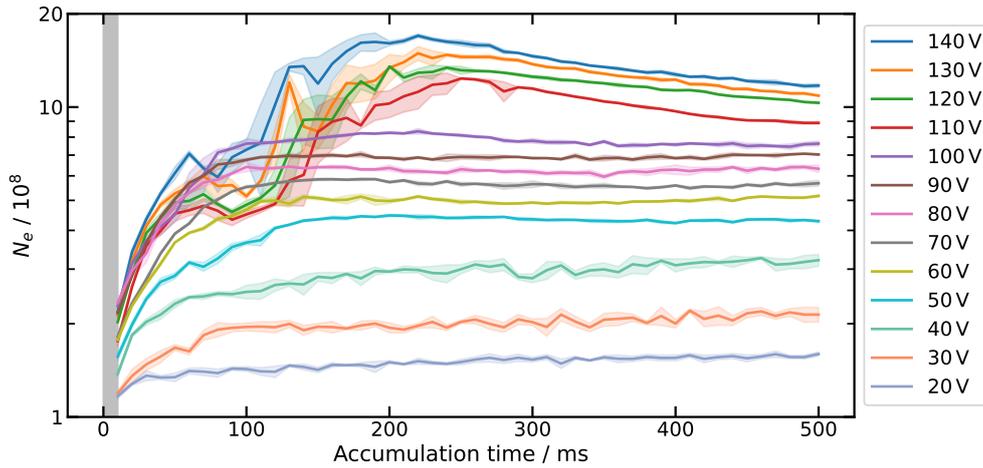
Figure 5.12: Total number of electrons trapped as a function of the time duration during which the particles were accumulated in the trap. The labels indicate the bias voltage applied to all the electrodes of the trap with the exception of the end-cap electrodes and it is approximately equal to the 'depth' of the electrostatic potential well that confines the electrons axially. The grey area marks the time taken to set up the bias voltages which were maintained fixed during the accumulation of electrons. The shaded areas indicate the standard deviation measured for each curve.

5.3.3　Space-charge accumulation

The theoretical treatment of a magnetically confined plasma places clear limits on the maximum space-charge that can be trapped with fixed magnitudes of the axial magnetic field and of the depth of the electrostatic potential well. In reality, the quality of the vacuum, the misalignment between the magnetic field and the axis of symmetry of the electrodes, and instabilities of the plasma can lead to much lower achievable particle densities. An evaluation of the maximum achievable number of electrons in the storage trap at Swansea was carried out through a study of how the trapping efficiency changed for several confinement voltages, magnetic fields, and the duration for which the electrons were continuously accumulated. Trapping the electrons via the static method (slow accumulation) resulted in a total number of confined electrons of one or two orders of magnitude larger than those achievable with the dynamic technique (rapid beam capture). However, the maximum voltage of 140 V to which the electrodes can be biased imposes a limit on the maximum density that can be confined.

　　The accumulation of electrons was studied through the capture-release sequence outlined in table 5.1, with no storage stage as the particles were ejected immediately after the accumulation time interval had finished. At the start of the sequence, all electrodes were set to a high negative voltage to block the beam. Subsequently, the electrons were captured by setting the upstream end-cap electrode (E20) to a bias voltage of $-5\,\mathrm{V}$, the downstream end-cap electrode (E1) to $-140\,\mathrm{V}$, and the rest of the electrodes (E19–E2) to $V_0 > 0$. Here, V_0 dictated the 'depth' of the axial electrostatic potential well in which electrons were trapped. The bias voltages on the electrodes were changed in small increments for an initial interval of $10\,\mathrm{ms}$. Afterwards, the electrostatic potential-well was kept unchanged for a certain period of time during which more electrons were accumulated. After the accumulation stage, electrons were dumped onto the detector to record the number of trapped electrons
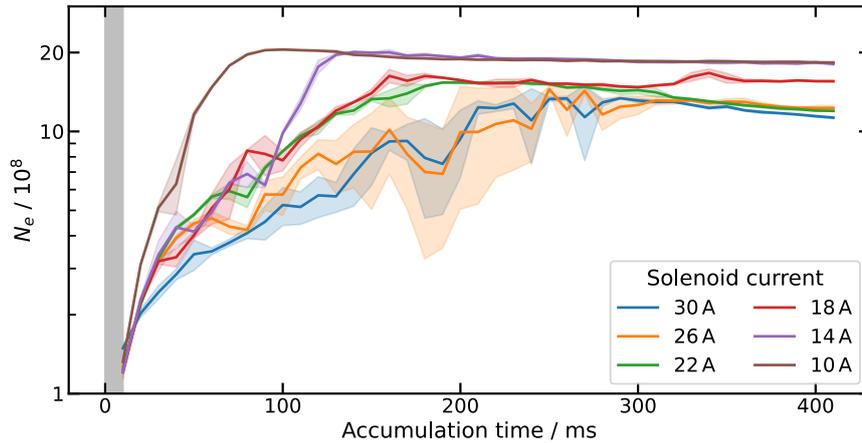
Figure 5.13: Total number of electrons trapped as a function of the time duration during which the particles were accumulated in the trap. The labels indicate the current flowing through the solenoid of the electron trap. The bias voltage applied to all the electrodes of the trap with the exception of the end-cap electrodes during the accumulation was 120 V. The grey area marks the time taken to set up the bias voltages which were maintained fixed during the accumulation of electrons. The shaded areas indicate the standard deviation measured for each curve.

for several values of the bias voltage V_0.

The number of electrons trapped for several bias voltages applied to the electrodes E19 to E2 is shown in fig. 5.12 as a function of the interval of time for which the particles were accumulated. A higher voltage on electrodes E19–E2 is equivalent to a 'deeper' longitudinal electrostatic potential well and, thus, leads to more electrons being trapped. The measurements show that there is an optimal interval of time during which the electrons should be accumulated in order to maximise the number of electrons that remain trapped. If electrons are injected into the trap beyond this optimal interval of time, the rate of loss of particles becomes comparable to the rate of accumulation, and the total number of electrons remains approximately constant or even starts to decrease. Furthermore, fig. 5.13 shows that the number of trapped electrons is also affected by the strength of the axial magnetic field of the trap. A higher magnetic field leads to a more compressed electron cloud and, thus, stronger space-charge repulsion, which may lead to a faster radial transport and particle loss.

The overall limit on the total number of electrons accumulated in the trap was found to be 2×10^9 particles. By the time the trap fills with a number of electrons close to this limit, the particle ensemble expands radially and mostly fills the entire volume of the trap. An evaluation of the corresponding peak electron densities was unachievable due to the damaged areas of the detector. The less sensitive areas led the transverse distributions recorded for the electron ensembles filling the entire trap to be incomplete.

The measurements described so far in this section were aimed at characterising the electron plasmas confined within the storage trap at Swansea. The rest of the section presents an investigation of the radial rotation of the electrons, i.e. the magnetron motion, comprising mostly of measurements of the magnetron frequency. Towards the end of the section, the measured trends are compared with predictions from computational models
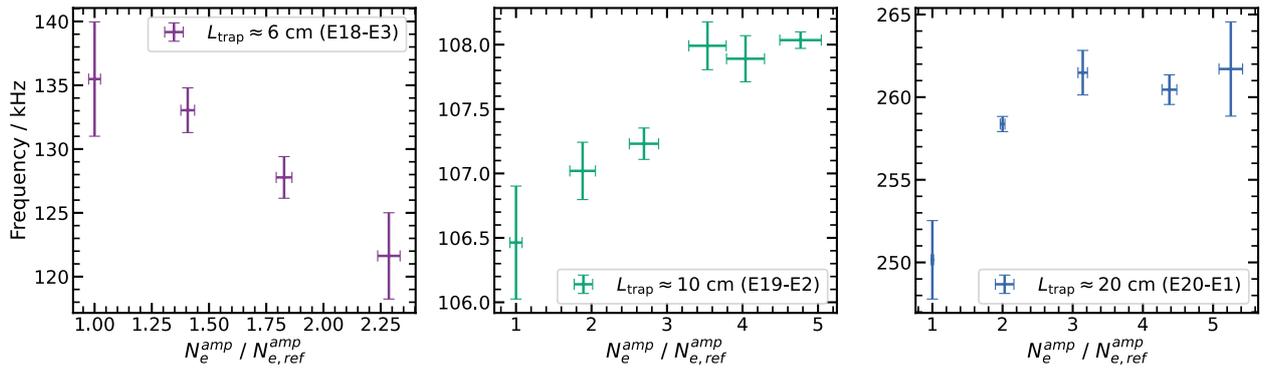
Figure 5.14: The change in the magnetron frequency with the total number of trapped electrons for three trap configurations with increasing length. The number of particles is normalised to a reference value chosen as the measurement with the lowest magnitude.

of the trap and simulations of the dynamics of the confined electrons.

5.3.4   Magnetron motion

A plasma column of charged particles that is radially offset from the centre of a Penning-Malmberg trap undergoes an azimuthal rotation around the axis of symmetry of the trap. When particles are released from the trap towards the imaging detector, the azimuthal position at which the particles hit the screen is correlated to the azimuthal position of the plasma column inside the trap. Thus, one can sample the transverse rotation of the plasma inside the trap by recording where the particles hit the screen at equally-spaced time intervals for which the plasma is allowed to rotate inside the trap before being ejected towards the detector. As a result, the frequency of the rotation, ω_m, of the particles inside the trap is determined from the frequency of the rotation of the spot where light is recorded on the phosphor screen.

Dependency of ω_m on the linear electron density

The transverse rotation of charged particles inside a Penning-Malmberg trap is dictated by two distinct mechanisms and can be classified as either magnetron or diocotron motion. The two types of motion can be distinguished by whether the motion of a single particle inside the trap is dominated by the electrostatic potential of the trap or by the space-charge of the entire particle ensemble (single-particle versus plasma regime). For example, the frequency of the magnetron motion is affected by the bounce frequency and, thus, by the length of the electrostatic potential according to eq. (5.9). Conversely, the diocotron motion is largely unaffected by the potential of the trap. Furthermore, as the diocotron motion is a direct consequence of the image charge induced in the conducting walls of the trap, the diocotron frequency is proportional to the linear density of trapped particles [235].

In order to identify which of the two types of motion was observed experimentally, the frequency of the transverse rotation was measured for three configurations of the trap leading to different plasma lengths and over a range of total electron number. The results shown in fig. 5.14 indicate that the transverse rotation of

Table 5.2: The experimental sequence and electrode voltages used in the study of the magnetron motion of trapped electrons. The example in this table corresponds to using electrodes E19 and E2 as end-caps for the trap. Equivalent sequences were used for the other two combinations of electrodes as end-caps (E20-E1, E18-E3). Δt is a time delay that determines at which phases during the magnetron rotation the electrons were released towards the imaging detector. The voltage bias supplied to the end-caps, V_{ec}, was varied between $-40\,\mathrm{V}$ and $-140\,\mathrm{V}$.

Phase	Duration	Voltages [V]						
		E20	E19	E18	E17–E4	E3	E2	E1
Reset	$100\,\mathrm{ms}$	0	-35	0	0	0	V_{ec}	0
Beam capture	$10\,\mathrm{\mu s}$	0	-20	0	0	0	V_{ec}	0
Store	$(10 + \Delta t)\,\mathrm{\mu s}$	0	V_{ec}	0	0	0	V_{ec}	0
Release	$10\,\mathrm{\mu s}$	0	V_{ec}	0	0	0	0	0

electrons observed experimentally corresponds to the magnetron motion. The measured frequency for each of the three axial extents of the trap show little variation even though the number of trapped particles and, hence, the linear electron density, change by at least a factor of two and up to a factor of five. Furthermore, the average frequency was observed to change between the three trap configurations which indicates that the transverse rotation is affected by the bounce frequency of the electrons inside the trap. As the linear density of electrons is expected to be approximately equal no matter how long the trapping volume was, no change in frequency is expected to the first order if the rotation corresponds to the diocotron motion. The conclusion that the observed rotation corresponds to the magnetron motion is also confirmed below from comparisons to simulations of the trapped electrons. Once the rotation had been identified as the magnetron motion, the change in the magnetron frequency was investigated over several strengths and configurations of the confining fields of the trap.

For the study of the magnetron motion of electrons inside the trap, the particles were captured via the dynamic technique. The stages of the experimental sequence are given in table 5.2. Hence, the resulting electron columns were thin and injected off-axis into the trap so that they could undergo a transverse rotation around the symmetry axis of the trap. However, the total number of electrons that could be captured via the dynamic procedure was limited by the electron beam current. By contrast, the static accumulation of electrons could have lead to larger number of trapped particles, but the resulting electron column was typically wide and centred such that no magnetron rotation was subsequently observed.

Figure 5.15 shows some of the average quantities of interest measured during several experimental runs with electrons undergoing magnetron motion. The electrons were trapped via the dynamic technique through which the electron beam from the source is rapidly captured inside the confinement volume. Each data point shows the average of 200 capture-store-release sequences for a fixed voltage applied to the end-cap electrodes of the trap. The radius of the transverse rotation of the electron column was observed to change with the end-cap voltage. Since the electrons were confined by directly capturing the beam, the resulting electron columns were thin, beam-like, with average full-width-at-half-maximum of approximately $0.7\,\mathrm{mm}$ inside the trap. The total

number of particles that were captured via the dynamic technique was in the range 0.5–2.0×10^6.

The average electron densities in fig. 5.15 were calculated from the measured values of the transverse size and total number of particles, assuming that the length of the plasma column corresponds to the distance between the centres of the end-cap electrodes. The resulting electron densities are of the order of $10^{13}\,\mathrm{m}^{-3}$. If we assume that the electrons were trapped purely by capturing the beam, the typical calculated electron density corresponds to a beam current of around $3\,\mu\mathrm{A}$. The decreasing trends observed in the electron density are caused by the assumption of a fixed length of the electron column. In reality, increasing the end-cap voltage changes the shape of the axial electrostatic potential which decreases the length of the region where electrons are confined. The step change in the total number of trapped electrons, and hence, in the electron density, around an end-cap voltage of $70\,\mathrm{V}$ appears in all three experimental runs which were performed a couple of hours apart. Together with the variation in the electron density between the three experimental runs, the step change in the total particle number shows that the dynamic trapping technique was not fast enough simply to capture part of the electron beam. Since the rise time of the voltage applied to the end-cap electrodes was limited to $3\,\mu\mathrm{s}$, other processes might have led to electrons being trapped in addition to capturing part of the beam. For example, the electrons that enter the trap earlier undergo additional magnetron rotations and axial bounces compared to the electrons that arrive just before the voltage on the end-cap electrode is large enough to completely stop further electrons from entering the trap.

The radius of the transverse rotation of the electrons was typically several times the FWHM of the electron column and could be varied by changing the vertical position of the electron source as indicated in fig. 5.16. Minimising the amplitude of the magnetron rotation can be used as a technique to find the vertical position of the source which ensures that electrons are injected close to the axis of symmetry of the trap. The fact that the minimum amplitude of rotation in fig. 5.16 is still significantly above zero suggests that the axial magnetic field was not perfectly aligned with the axis of symmetry of the electrodes. Figure 5.16 also shows a small variation of a few percent of the magnetron frequency with the amplitude of the rotation.

Higher electron densities with modified source

The magnetron frequencies that were reported so far in this work were measured with the unmodified electron source. Removing the Wehnelt cylinder from the source led to a larger electron column in the trap and a higher number of trapped electrons. An additional experimental run was dedicated to measuring the magnetron frequencies with a higher total number of confined electrons. The aim was to verify whether the space-charge of the electrons was high enough for the diocotron motion (collective effect) to be observed instead of the magnetron motion (single-particle regime). A similar batch of sequences was run as before to measure the change in the transverse rotation frequency of the electrons with increasing beam current and, hence, the number of electrons being confined. The results are presented in fig. 5.17 and 5.18. Even though the total number of electrons increases by more than a factor of eight in these measurements (fig. 5.17), the rotation frequency did not show a variation of more than 10%. Hence, the measured transverse rotation is a consequence of the
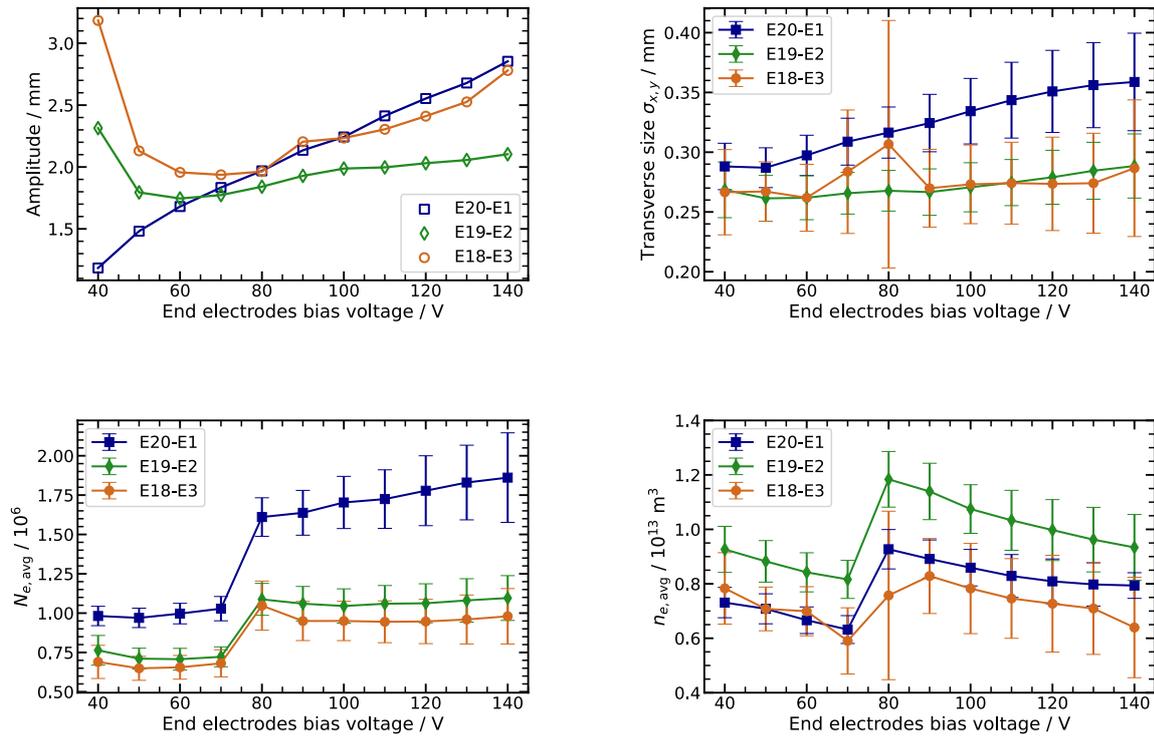
Figure 5.15: Average quantities measured for individual experimental runs for the study of the magnetron motion of electrons as a function of the trapping voltage applied to the end-cap electrodes: the amplitude of the transverse rotation of the electrons in the trap, the transverse size of the electron column in the trap, the total number of trapped electrons ($N_{e,\mathrm{avg}}$), average electron density ($n_{e,\mathrm{avg}}$). The labels indicate separate experimental runs, where E##-E# specifies which electrodes were used as end-caps in the particle trap.
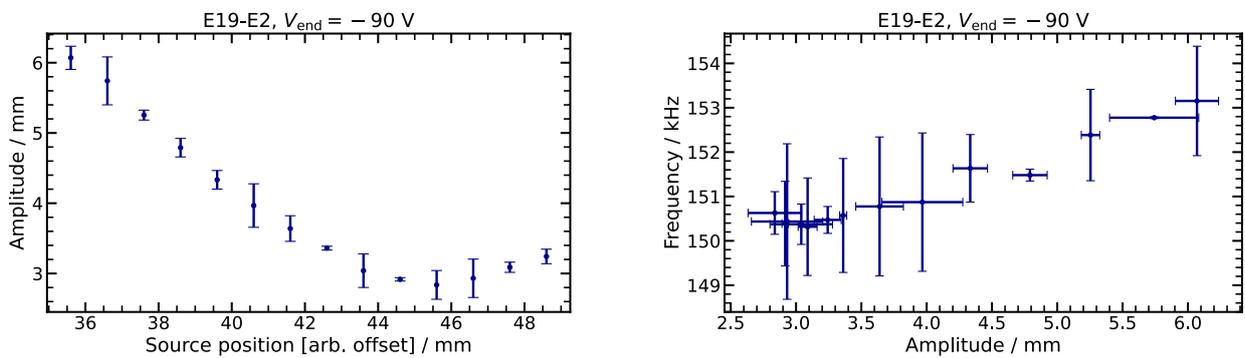
Figure 5.16: The variation of the amplitude of the magnetron motion of trapped electrons with changes in the vertical position of the particle source (left). The variation of the magnetron frequency over the range of measured amplitudes (right).
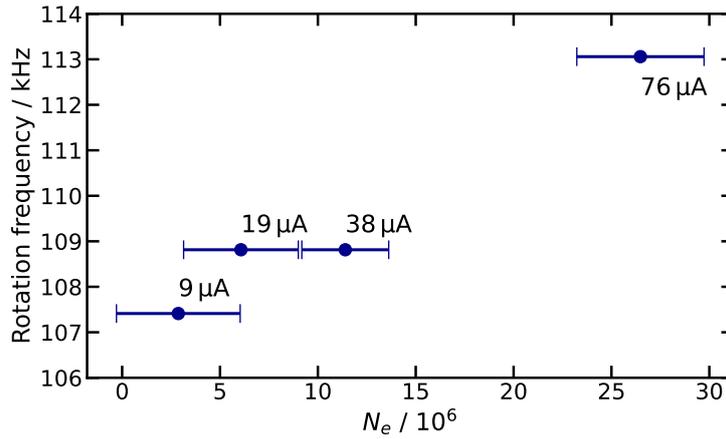
Figure 5.17: The frequency of the transverse rotation of electrons measured for several filament emission currents. The trapping sequence was identical. The horizontal axis shows the average number of electrons trapped as measured for each emission current indicated on the plot. The transverse size and the electron density of the plasma columns during these experimental runs are shown in fig. 5.18.
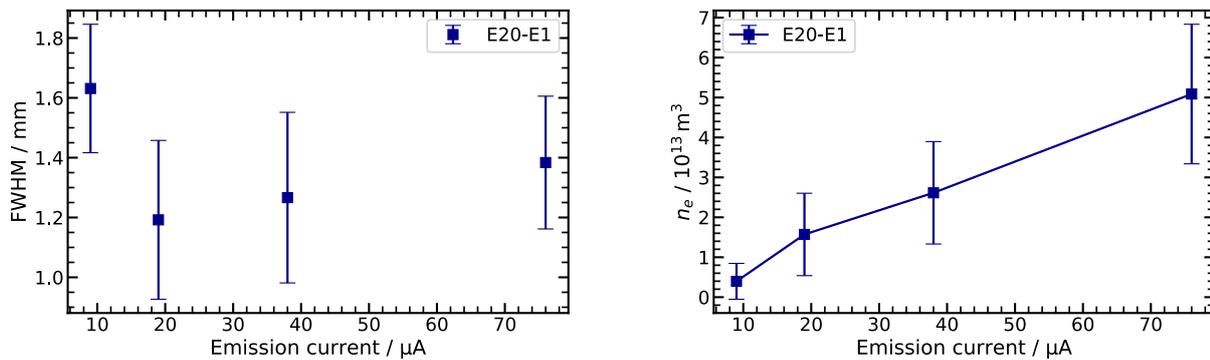
Figure 5.18: Transverse size (left) and particle density (right) of the electron clouds confined during the measurements of the magnetron frequencies in fig. 5.17 with the modified electron source. The densities show on the right represent lower limits since the length of the electron column was assumed to be $24\,\mathrm{cm}$ (the distance between electrodes E20 and E1).

magnetron motion of the electrons inside the trapping potential. Moreover, the space-charge of the electrons was not high enough for collective effects, such as the diocotron motion, to be measured, even with the use of the modified source.

In the measurements performed with the modified source and summarised in fig. 5.18, the peak electron density was estimated to be $\approx 5 \times 10^{13}\,\mathrm{m}^{-3}$ ($N_e/L = 1.1 \times 10^8\,\mathrm{m}^{-1}$) for the experimental run with an emission current of $76\,\mu\mathrm{A}$. This is the overall upper limit of the electron densities that were achieved throughout all the experimental work presented in this chapter. The large error bars in fig. 5.18 are caused by the large variations in the transverse size of the light intensity profile created by the electrons on the phosphor screen. When the electrons hit the screen in areas with more screen damage, significant deviations from a Gaussian intensity profile were observed.
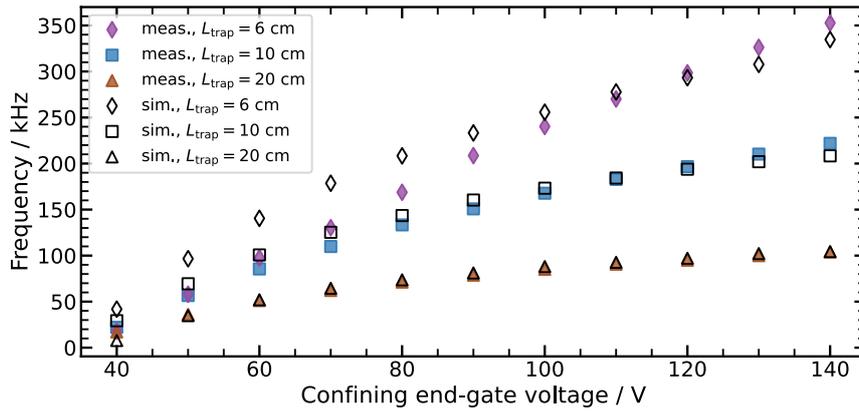
Figure 5.19: The magnetron frequency of the trapped electrons as a function of the voltage applied to the end-caps of the particle trap. The full markers indicate the results from measurements while the empty markers show the results of simulations. Each curve corresponds to a specific pair of electrodes used as end-caps of the trap and is labelled by the approximate length of the stored-electron column.

Dependency of ω_m on the strength of the trapping fields

As the transverse rotation of the electron columns had been identified as a result of the magnetron component of the motion of each single particle inside the external fields of the trap, two additional series of measurements were performed to investigate how the magnetron frequency varies with the strength of the external fields. These measurements were also identified as a good reference against which the particle-in-cell simulations in VSim could be verified because of the relatively short time scale of the magnetron period in the range of $2\,\text{ms}$ to $200\,\text{ms}$. With an appropriate choice of the simulation volume, mesh size, and time step, a PIC simulation can cover the evolution of a low-density electron plasma over such a time scale with moderate requirements for CPU time.

In the first series of measurements, the magnetron rotation was recorded by scanning the bias voltage applied to the electrodes chosen as end-caps of the trap starting from $-40\,\text{V}$ and down to $-140\,\text{V}$. This procedure was repeated for three pairs of electrodes used as end-caps such that the effective length of the axial potential well was approximately 6, 10, and $20\,\text{cm}$, respectively. Figure 5.19 shows the trends obtained by extracting the magnetron frequency from the measurements resulted in the scans of the end-cap bias voltages. Each data point corresponds to a different shape and length of the axial electrostatic potential well and, thus, to a different axial bounce frequency of an electron. The magnitude of the current supplied to the solenoid of the trap was kept fixed throughout these measurements. As a result, the variation of the magnetron frequency is a consequence of the change in the axial frequency of the electrons bouncing from one end-cap electrode to the other in distinct potential wells.

In an ideal Penning trap, a higher voltage applied to the end caps leads to a higher bounce frequency and, thus, to a higher magnetron frequency. (see eqs. (5.7) and (5.9)). This general trend is also present in fig. 5.19. Furthermore, shorter traps result in higher bounce frequencies for electrons with the same kinetic energy. Figure 5.19 also includes the magnetron frequencies extracted from simulations of the motion of
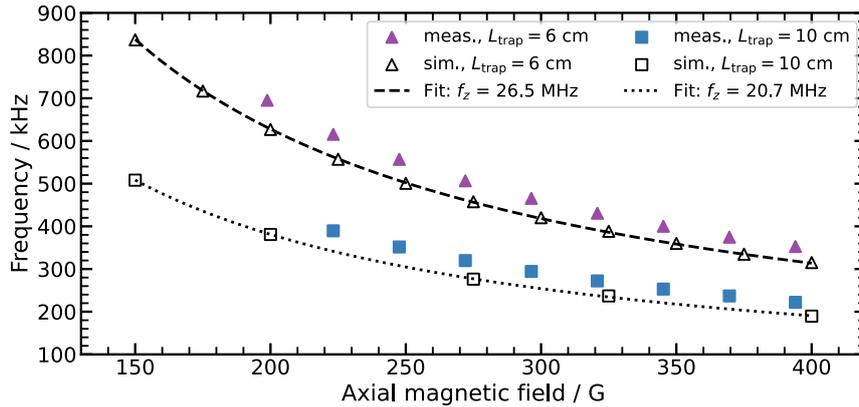
Figure 5.20: The magnetron frequency of the trapped electrons as a function of the magnetic flux density generated by the solenoid of the trap. The full markers indicate the results from measurements while the empty markers show the results of simulations. Each curve corresponds to a specific pair of electrodes used as end-caps of the trap and is labelled by the approximate length of the stored-electron column.

electrons inside the apparatus. A more detailed description of the simulation model is provided in the final part of this section. The measured data and the results of the simulation agree well for the $6\,\text{cm}$ and $10\,\text{cm}$ traps. For the $6\,\text{cm}$ case, there is a clear difference in the trend of the measured data compared to the prediction from the simulations which is more pronounced for the lower end-cap voltages. The discrepancy was attributed to the energy spread of the electrons during the measurements. In the simulations, the electrons have no energy spread and a total energy of $25\,\text{eV}$. It was shown before that the electrons captured during the experiments had a large energy spread. Hence, the average bounce frequency of the electrons in the case of the measurements is believed not to be fully matched by the bounce frequency of the electrons in the simulations.

In a second series of measurements of the magnetron rotation, the data was recorded by scanning the current supplied to the solenoid magnet of the trap from $14\,\text{A}$ to $30\,\text{A}$. This procedure was repeated for two pairs of electrodes used as end-caps to fix the length of the trapping volume to $6\,\text{cm}$ and $10\,\text{cm}$, respectively. Figure 5.20 shows that lower magnetron frequencies were measured at stronger magnetic flux densities according to a monotonously decreasing trend. This decrease is expected since in an ideal Penning trap, a lower magnetic flux density, and thus a lower cyclotron frequency, leads to a lower magnetron frequency (see eq. (5.9)).

Figure 5.20 includes the predictions from simulations of the motion of electrons in a model of the trap. The magnetron frequencies extracted from the simulations show good agreement with the measurements. Moreover, the values obtained from the simulations are fitted with eq. (5.9) with the axial bounce frequency as a free parameter. There is a good agreement between the theoretical model of the magnetron motion in an ideal Penning trap and the magnetron frequencies reproduced with the simulation model of the actual trap. The fact that the measured values lie slightly above the predictions of the simulations may be a result of a mismatch between the bounce frequency of the electrons in the simulations and in the experiments which is equivalent to a mismatch in the energy distributions of the electrons in the two cases. According the to fitted curves, the average energy of the electrons during the measurements was higher than the value chosen for the simulations.

Table 5.3: The modelling parameters of the simulation of the magnetron motion of trapped electrons. The magnitudes were chosen to match the measurements where data was available.

Model parameters	Magnitude
Magnetic field	15 to 40 mT
Total length of the grounded trap electrodes	63 mm, 101 mm, 201 mm
Inner electrode radius	20.5 mm
End electrodes potential	-40 to -140 V
Initial plasma transverse size	$\sigma_{x,y} = 0.3$ mm
Initial plasma transverse offset	1 mm
Initial plasma length	38 to 210 mm
Initial plasma density	3×10^{13} m^{-3}
Initial space charge on z axis	0.086 V
Total electron energy	25 eV

An alternative source of the discrepancy between the measured data points and the predictions of numerical model may be the inaccurate calculation of the magnetic flux density inside the trap. The values on the horizontal axis of fig. 5.20 were obtained by converting the values of the current supplied to the solenoid into magnetic flux density with the aid of a finite-element model of the magnet. Any deviation of the strength of the magnetic field produced by the solenoid compared to the strength calculated from the model introduces a shift of the measured data points along the horizontal axis of fig. 5.20. The final part of this section provides more details on the simulations to describe how the experimental measurements were reproduced and to highlight the limitations of the numerical model.

5.3.5 Simulation model

The electron trap employed during the measurements at Swansea University was modelled in VSim [233] and the motion of the trapped electrons was simulated with the particle-in-cell (PIC) method via the computational engine of VSim [243]. The parameters of the electron clouds that were measured during the experimental runs and the specifications of the storage trap were the main inputs to model the motion of electrons that was observed experimentally. The most important model parameters are given in table 5.3. In addition to these parameters, the simulations were set up to model the motion of individual electrons (see table 5.4). For example, the time step was chosen to be significantly smaller than the cyclotron period in order to resolve fully the rotation of the electrons around the magnetic field lines. The results of the PIC simulations are presented above in fig. 5.19 and 5.20 in conjunction with the magnetron frequencies measured experimentally. The most important aspects of the simulation model are discussed below.

The plasma simulations in VSim are based on the particle-in-cell (PIC) method. In the case of the simulations for the magnetron motion, a discrete number of electrons are combined in the so-called macro-particles. The

Table 5.4: Values chosen for the most important parameters of the simulation of the magnetron motion of trapped electrons.

Simulation parameters	Magnitude
Macroparticles per cell	64
Mesh size	$\Delta x = \Delta y = \Delta z = 0.2\,\text{mm}$
Max. total number of cells	$25 \times 25 \times 1395$
Time step	$\approx 80 \times \tau_c$
	$\approx 3 \times 1/\left(v_{\text{e.max}} \cdot DLI\right)$

$\tau_c = \left(2\pi m_e\right)/\left(eB_{\text{max}}\right)$ is the minimum cyclotron period.

$DLI = \sqrt{1/\Delta x^2 + 1/\Delta y^2 + 1/\Delta z^2}$ [244]
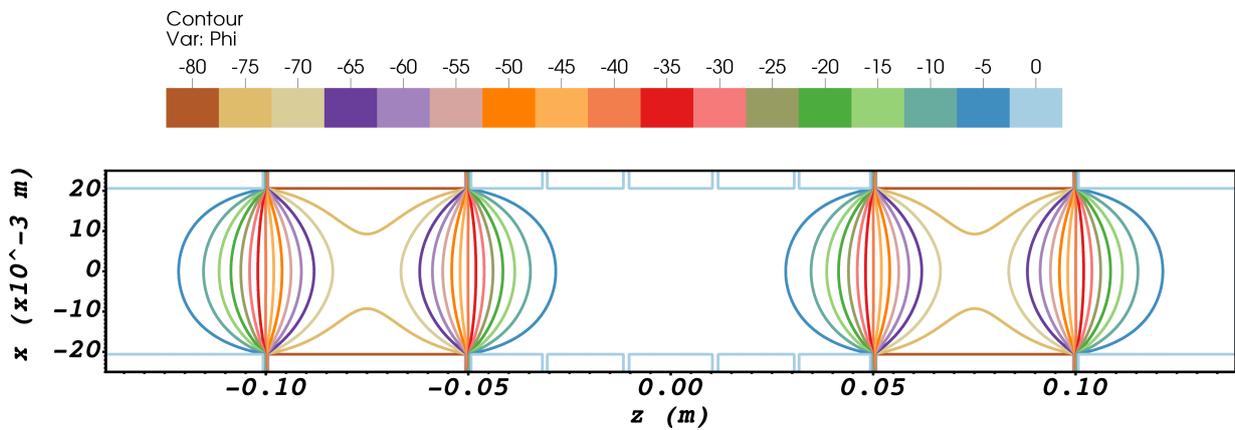
Figure 5.21: Equipotential contours of the electrostatic potential produced by electrodes E19–E2 for a bias of $-80\,\text{V}$ as obtained with VSim (longitudinal mid-plane).

Lorentz equation of motion is employed to advance the macro-particles in time. In the canonical PIC method, the charge density is calculated after each time step via a linear interpolation of the macro-particles position onto a mesh. However, in order to increase the computational efficiency, the simulations discussed here employed a different method of 'prescribed fields', which denotes the fact that the electrostatic and magnetic fields are kept constant throughout the simulation and the self space-charge of the electrons is neglected.

The method of 'prescribed fields' is suitable here because of the low density of the electrons that was measured during the experimental runs and the small transverse size of the electron column. The Debye length for the electron columns trapped in the measurements of the magnetron rotation is in the range $5\,\text{mm}$–$10\,\text{mm}$ which is significantly larger than the transverse size of the electron columns that were trapped. Thus, the electrons were confined in a 'single-particle' regime rather than a multi-particle, plasma regime. As a consequence, the motion of individual electrons is mostly governed by the external trapping fields and not by the interaction with the rest of the electrons. Furthermore, the space-charge of the electrons was below 0.3% of the trapping potential which means that the confining electrostatic well is largely unaffected by the presence of the charged electron column.

The main advantage of the 'prescribed fields' method over the canonical PIC procedure is that the simulation
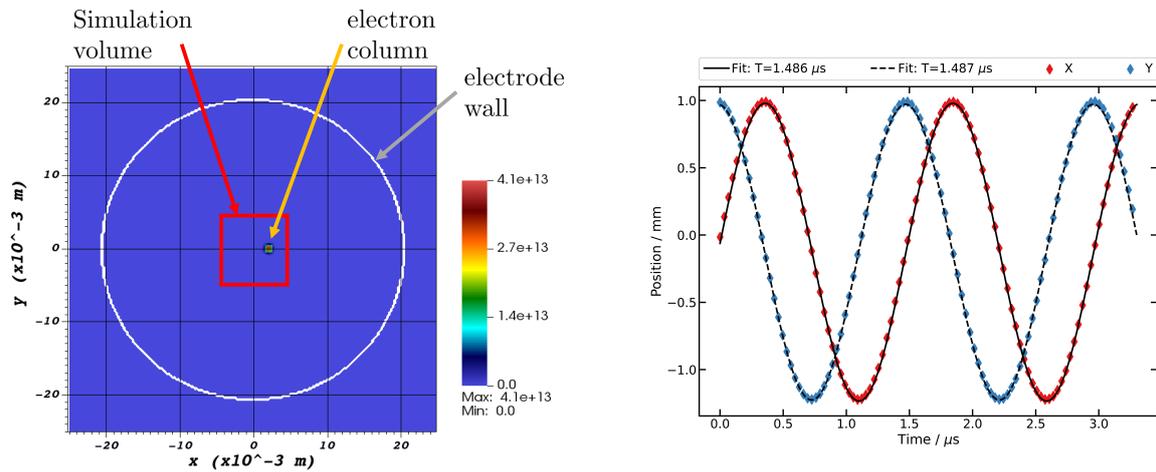
Figure 5.22: **Left:** Transverse cross-section through the simulation model for the magnetron motion of trapped electrons. The 'prescribed fields' method allows a reduction of the simulation volume from the full area in blue to the red square. The colour bar on the right shows the electron density. **Right:** Example of the position of the electron column versus time during a simulation.

domain could be significantly reduced. In the transverse plane, the size of the electron cloud and the radius of the magnetron rotation were both of the order of $1\,\mathrm{mm}$. However, the conducting walls of the electrodes are at a radius of $20.5\,\mathrm{mm}$ from the axis of the trap. Hence, most of the simulation domain is not populated by macro-particles if the conducting boundaries of the electrodes have to be included in the simulation. With the full PIC method, the Poisson equation is solved after each time step and, hence, the conducting boundaries of the electrodes have to be part of the simulation. At the same time, the mesh size has to be kept small in order to resolve the radial profile of the electron column. As a result, the total number of cells becomes large and the simulation requires an increased computational capacity. With the 'prescribed fields' procedure, the external electrostatic field produced by the electrodes is calculated once at the beginning of the simulation in the absence of electrons. An example of the electrostatic potential calculated with VSim is shown in fig. 5.21. The magnetic field is imported in addition to the electrostatic potential well and both are kept constant for the rest of the simulation. Finally, the simulation boundaries are chosen to tightly include the volume in which the electrons move, thus excluding any mesh cells that macro-particles never reach. Figure 5.22 illustrates the relative size of the electrodes, electron column and simulation volumes.

The configuration and dimension of the electron trap in the simulation model as well as the magnitudes of the electrostatic potentials are consistent with the experimental setup. However, the magnitude of the axial magnetic field was separately calculated with a model of the solenoid. The strength of the magnetic field was varied during the experimental runs by changing the current through the magnet. A finite-element magnetostatic model of the solenoid was created to obtain the magnetic field distribution inside the solenoid for a set of currents (see fig. E.8). The distribution of the magnetic field was subsequently imported into the PIC simulation.

Simulations were carried out on discrete square mesh cells in Cartesian xyz coordinates. For the initial calculation of the external electrostatic potential created by the electrodes, Dirichlet conducting boundary condi-

tions were used on the internal surfaces of the electrodes. The mesh size was chosen in order to resolve the radial profile of the electron column. Similarly, the magnitude of the time step was chosen to resolve the cyclotron motion of the electrons around the magnetic field lines.

The initial electron distribution was chosen to match the profiles measured in the experimental runs. For the radial profile, the density was initialised according to a radially symmetric Gaussian distribution, while in the longitudinal direction, the electron density corresponded to a uniform column with hard edges at the two ends. The initial length of the electron column was changed to correspond to the voltage applied to the end electrodes. In order to do that, the length of the electron column was determined from the maximum distance from the mid-plane of the trap that electrons can reach in the electrostatic potential created by the electrodes. Since the electrons were rapidly trapped as a beam during the experimental runs, the longitudinal energy of the electrons was set to $25\,\mathrm{eV}$ in the simulation. During the initialisation of the particles, the longitudinal energy of the electron was set accordingly to take into account the electrostatic potential well. Thus, electrons closer to the ends of the trapping volume were initialised with a lower energy to account for the beam propagation through the potential well. Finally, in the two transverse directions, the electrons were initialised with a Maxwell-Boltzmann velocity distribution corresponding to a temperature of $1\,\mathrm{eV}$. An example of an initial distribution of electrons in spatial coordinates and velocity space is shown in fig. 5.23.

5.4 Discussion

A series of measurements of low-density non-neutral electron plasmas were performed on the positron beam-line at Swansea University, in particular by using one of the available Penning-Malmberg traps. The experiments were facilitated by the existing particle detector assembly that provided destructive diagnostics of the trapped electrons, and by the control and data acquisition systems. Prior to this work, extensive studies have been done on the beam-line at Swansea Unviersity for many years in order to trap, store, and characterise positron clouds, as well as to develop the necessary experimental apparatus and techniques. An investigation of non-destructive plasma diagnostic techniques reported the initial experiments with electrons confined within the same beam-line [241]. The work presented in [241] includes the early measurement of the magnetron frequency of electrons in the single particle regime (low density electron clouds) in one of the positron traps at Swansea University. Most of the measurements described in this chapter represent the first set of attempts to trap or accumulate electron clouds with a high enough density for the electrons to form a well-defined plasma. The final objective was to observe the transition from the single-particle regime to the plasma regime from measurements of the transition from the magnetron motion of the electrons to the diocotron motion.

Ultimately, it was found in this work that the electron densities achieved were too low for the transition to the plasma regime to be observed. However, the measurements investigated the limitations of the storage trap by characterising the electron plasmas confined with various configurations of the trap. While the capabilities of the beam-line to trap and store positrons were well known, a more detailed study of the electron plasmas
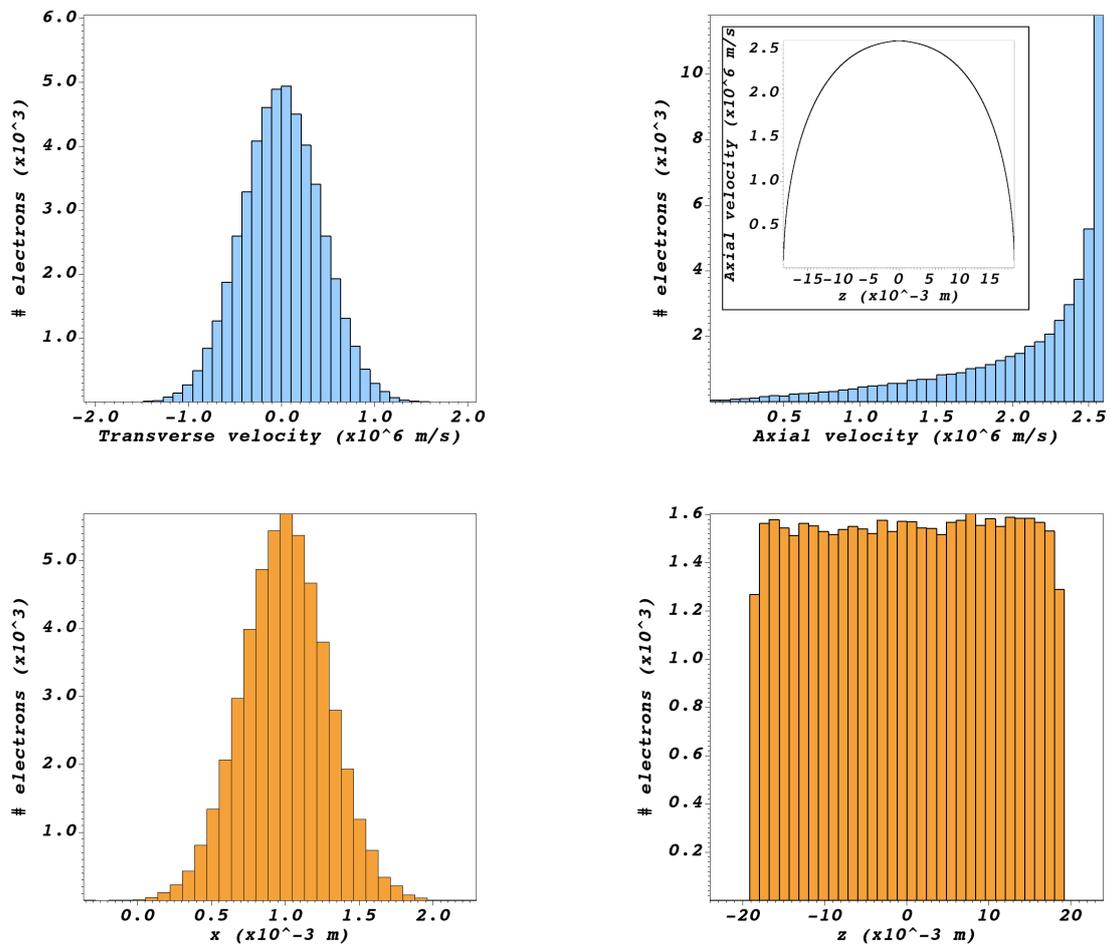
Figure 5.23: Example of an initial electron distribution in the transverse (x) and longitudinal (z) directions used for the simulation of the magnetron motion. The distribution corresponds to a trap configuration where electrodes E18,E3 are used as end-caps and the applied bias is $140\,\mathrm{V}$. The transverse density profile is Gaussian while the longitudinal one is uniform with a hard edge. In the velocity space, the transverse component is sampled from a Maxwell-Boltzmann distribution, while the axial velocity is matched to the electrostatic potential well such that electrons have a total longitudinal energy of $25\,\mathrm{eV}$.

Table 5.5: Summary of the electron plasma parameters measured during the experiments on the positron beam-line at Swansea University.

Measured plasma parameter	Magnitude
Length	6–20 cm
Radial FWHM	0.6–30 mm
Total number of electrons	$\lesssim 2 \times 10^9$
Electron density	$\lesssim 5 \times 10^{13} \mathrm{m}^{-3}$
Lifetime	0.6–9 s

that can be achieved within the same bealine has not been done before. Furthermore, the previously existing electron source was modified to provide a larger electron-beam current and the capabilities of the electron trap were investigated for the first time with the use of the modified source.

A secondary goal of the measurements described in this chapter was to compare suitable data from the experiments with predictions from particle-in-cell simulations of the trap and of the motion of electrons carried out with the VSim software. VSim was used in this work to model the electron plasma produced by a Gabor lens and to reproduce some experimental results obtained with the lens. Hence, the software was identified as a simulation tool with a great potential to serve at making predictions on the capabilities of an improved version of the lens that had been tested. Thus, some of the data acquired with the use of the electron trap at Swansea University was planned to serve a validation of the predictions obtained with VSim. To the author's knowledge, VSim and the associated computational engine Vorpal have not been tested before against laboratory measurements in the case of modelling an electron Penning-Malmberg trap.

The first class of measurements yielded a better understanding of the plasma parameters achievable with the particle trap at Swansea as summarised in table 5.5. **The maximum electron density achieved was approximately** $5 \times 10^{13} \mathrm{m}^{-3}$ **and the maximum number of electrons** 3×10^7**.** The maximum number of trapped electrons was achieved through continuous accumulation of particles inside the axial potential well for periods of approximately 200 ms, after which the rate of particle loss becomes equal to that of particle accumulation.

Two conclusions can be drawn from the first class of measurements. Firstly, the **maximum number of trapped electrons was dictated by the limited beam current** and diameter generated by the hot filament. Hence, the development of the electron source must be part of the R&D work towards a high-density electron plasma lens. Secondly, the **active reduction of the outward radial transport of electrons is needed** to achieve quasi-steady central plasma densities for periods of time longer than about 100 ms. During these measurements, it was also evident that the main source of uncertainties and sometimes the main limitation in measuring the profiles of large electron clouds was the damaged MCP/P-screen assembly.

The second class of measurements investigated the average transverse rotation of the electrons inside the plasma trap. **The rotation was identified as the magnetron component of the electron motion** as the frequencies measured show little variation with changes in the linear density of trapped electrons. Moreover, the

rotation frequency was shown to depend on the strength of the external fields of the trap in a similar manner to the expectations from the theoretical model of an ideal Penning trap. In conclusion, the electron densities achieved within the trap at Swansea were too low for the electrons to behave as a non-neutral magnetically-confined plasma. **The predictions from the numerical simulations of the trap and the motion of the electrons are in good agreement with the measured values for the magnetron frequency in the range 20 kHz–700 kHz** considering the wide energy spread of the electrons during the experiments and the inexact knowledge of the magnitude of the magnetic flux density inside the trap.

Finally, a remark should be made regarding the field solver used for the simulations of the motion of electrons inside the trap. The electric and magnetic fields were kept constant in the simulation at every time step in order to achieve a moderate CPU time for each simulation. The fields were identical to the external fields generated by the trap. As a consequence, the space-charge of the electrons was not included in the field solver as the motion of individual electrons during the experiments was mostly dictated by the external fields at the low plasma densities studied here. This implies that the results presented above should be treated as a verification of the particle 'pusher' that is part of the VSim particle-in-cell solver and of the field solver in the absence of particles. In a full particle-in-cell simulation, the space-charge of the macro-particles are included in the calculation of the fields at the end of each time step. A further series of measurements with plasmas at higher densities is required to benchmark the full particle-in-cell method implemented in VSim.

In conclusion, the positron/electron beam-line, and the associated Penning-Malmberg traps, at Swansea University are a versatile setup that can serve dedicated measurements on well-confined, low-density electron clouds. Simulations of the magnetron motion of the electrons inside one of the traps of the beam-line showed that certain aspects of the experiment can be quantitatively mirrored by particle-in-cell simulations. Since the electron densities achieved during the experiments were comparatively low, the agreement between the simulations and the measurements in terms of the magnetron frequency provides a validation of the electrostatic solver and the correct modeling of the single-particle motion in VSim. The validation of the complete PIC method in VSim requires measurements of certain plasma parameters (e.g. the diocotron frequency) at higher electron densities where the particle ensemble can be considered a well-defined plasma for which long-range electron-electron interactions play a role in the behaviour of the electron cloud. Therefore, a new, stand-alone test-bench is required to perform experiments to investigate plasma parameters of interest in a regime that is more similar to that of the plasma lenses envisioned for LhARA.

The following chapter will describe a preliminary design of a normal-conducting (NC) solenoid magnet which could act as an alternative to the electron plasma lens planned for LhARA. In presenting the design, the requirements for strong focusing and large aperture will be shown to impose lower limits on the size and power-consumption of the magnet. After that, chapter 7 will return to beam focusing by plasma lenses and will evaluate the performance of the NC solenoids compared to that of the plasma lenses in the Stage 1 of LhARA.

Chapter 6

Design of a normal-conducting solenoid for LhARA

Solenoid magnets could present an alternative to plasma lenses to provide the strong focusing required in LhARA to capture and transport the protons and ions from the laser-driven source and reduce the beam energy spread. The risk mitigation strategy, which is part of the development of the conceptual design of LhARA, includes the design of solenoids required to construct an alternative particle-capture system [10]. Conventional solenoids can provide the focusing strength needed to replace the plasma lenses in the case of LhARA. However, the extension of the laser-hybrid technique to deliver proton and ion beams for clinical use would require high-field solenoids. In order to investigate the parameters of a realistic focusing magnet for LhARA, a design study was carried out, which resulted in three preliminary designs for a normal-conducting solenoid. The normal-conducting technology was chosen as the first option to be explored due to its extensive use in accelerator magnets and the relative simplicity of running a magnet close to room temperature. Moreover, the designs for the solenoids provide several figures of merit, such as power consumption and size, which can be compared to those of the equivalent plasma lenses.

This chapter begins with an overview of the requirements of a solenoid that is suitable for the LhARA beam-lines.(section 6.1.1). Then, the choice of normal-conducting technology for a first consideration of the LhARA solenoids is discussed (section 6.1.2). Following a description of the principles of magnet design (section 6.1.3), preliminary designs are presented for three solenoids which are planned to operate at three distinct currents (section 6.1.4). Lastly, the performance of the optimal designs according to numerical simulations is reported (section 6.1.5) and the parameters of the corresponding solenoids are discussed (section 6.2).

6.1 Normal-conducting solenoids for LhARA–properties and design

Solenoids are typically used in the low energy section of accelerators [125] to provide transverse focusing for charged particle beams or to capture the particles from the source. In addition, solenoids are a main component of non-neutral plasma traps and particle detectors. The quasi-uniform magnetic field produced by a solenoid is

```
┌──────────┐    ┌─────────────┐    ┌──────────────┐    ┌────────────┐    ┌──────────────┐
│Input data│───▶│ Analytical  │──▶│ Numerical 2D/3D│──▶│ Mechanical │──▶│  Drawings &  │
│          │    │   design    │    │  simulations  │    │   design   │    │specifications│
└──────────┘    └─────────────┘    └──────────────┘    └────────────┘    └──────────────┘
```
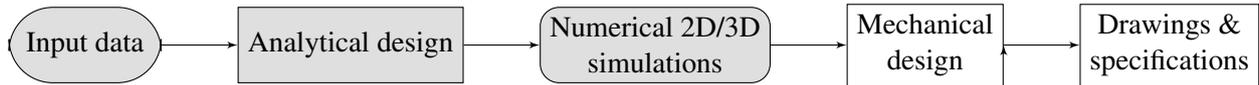
Figure 6.1: The main stages in the electromagnetic design of normal-conducting magnets for accelerators. The first three stages are the ones required for the conceptual design of LhARA and are addressed in this chapter.

of interest for confining charged particles to fixed volumes for long periods of time or for guiding the particles along curved trajectories in high-energy particle physics detectors. However, in a long solenoid, charged-particle beams are focused due to their passage from a field-free region outside of the magnet, through the uniform part of the solenoidal field and, finally, outside of the field region. It is the variation of the radial and axial components of the magnetic field at the two ends of the solenoid that results in an overall focusing of the beam. For a more detailed description of the single-particle motion in the magnetic field of a solenoid and an explanation of how the magnet acts on a beam to produce the focusing effect, the interested reader is referred to appendix B.2.

The goal of the design and construction stages of a solenoid is to produce a magnet that performs reliably according to the specifications with sufficient safety margins, at the lowest cost, and for the required operational lifetime. The electromagnetic design is an iterative process that requires several optimisations during each of its stages which are summarised in fig. 6.1. Once a design has been produced, the construction phase begins with prototyping, followed by series production and tests, and, finally, installation and commissioning. In the case of a solenoid, which is not an iron-dominated magnet, most of the effort goes in the design of the coil. The power requirements of the coil are typically used to compute a suitable cooling circuit which, then, leads to the selection of the conductor. Although general methodologies for designing normal-conducting coils and the associated cooling circuit exist (see, for example, [245] and [246]), magnet design requires several iterations and educated choices of parameters to generate an optimised solution. Finally, a cost estimate and optimisation should also feed into the design of the magnet from the initial stages.

6.1.1 Magnet requirements

Before starting the design of a magnet, all input parameters, requirements, constraints, and interfaces with other systems of the accelerator have to be known and well understood. Based on the most recent beam-optics design for LhARA Stage 1 and a re-optimisation of the field values after tracking of a beam generated from 2D PIC simulations [247], the solenoids required for the capture section of LhARA must produce an integrated field

$$\int B(z)dz \leq 1.233 \, \text{Tm}. \tag{6.1}$$

According to the beam-optics studies, the solenoid closest to the particle source requires the largest integrated field since it has to produce a field of $1.4387 \, \text{T}$ with an effective length of $0.857 \, \text{m}$. The total length available for one solenoid is $1.157 \, \text{m}$ which coincides with the flange-to-flange distance of the equivalent Gabor lens. The goal of the work presented here is to design a solenoid that has a length of $1 \, \text{m}$ and generates a magnetic field $B_0 = 1.3 \, \text{T}$.
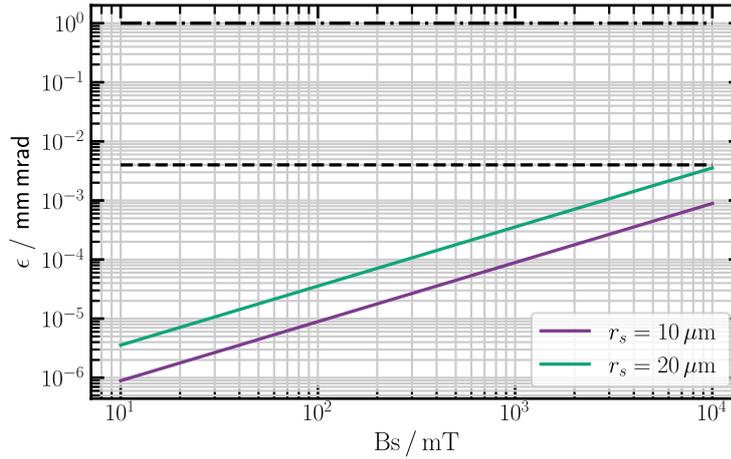
Figure 6.2: The magnetisation emittance ϵ obtained from eq. 6.3 as a function of the residual magnetic field B_s at a $15\,\mathrm{MeV}$ proton source with radius r_s. An example of the upper limit of the transverse emittance measured for a multi-MeV laser-driven proton beam [248] is also shown (dashed line). The approximate ideal beam emittance at the end station of LhARA Stage 1 [247] is further indicated (dash-dotted line). By comparison, the magnetisation emittance plays a negligible role in diverging the beam compared to space-charge forces and the beam transport up to the end-station.

The aperture of the magnet is mainly dictated by the maximum expected beam size in the second and the third solenoid in Stage 1 of LhARA, which is $\sigma_{x,y} \approx 15\,\mathrm{mm}$ [9]. Allowing some space for the beam pipe and an optional coil former or mesh, the inner diameter of the coil was chosen to be $76\,\mathrm{mm}$. As the size of the beam is comparable to the aperture of the solenoid, the 'good field region' is considered to be the full internal volume of the coil. Furthermore, since the requirements on the field quality were not generated for the initial design of LhARA, the solenoid discussed here should achieve the typical homogeneity value for accelerator magnets

$$\frac{\Delta \left(\int Bdz \right)}{\int Bdz} \leq 0.01\%. \tag{6.2}$$

Here, $\Delta \left(\int Bdz \right)$ is the difference between the actual integrated field and the ideal integrated flux density $\int Bdz$.

Lastly, the solenoid is planned to be operated in a continuous mode. Fast-tuning of the field strength is desirable in under $0.1\,\mathrm{s}$ to allow the delivery of different energies from consecutive pulses from the laser. In terms of the mechanical dimensions of the coil and the iron cover, no transverse space limitations have been identified. Although the first solenoid has to be placed close to the target, all the optical elements and the direction of incidence of the laser are expected to constrain the space upstream of the target and not downstream. It is important to note, that due to the layout of the capture section of the LhARA beam-line, there is limited space at the two ends of the solenoids to handle the electric and hydraulic connections.

Magnetisation emittance

A smaller distance between the laser-driven source and the solenoid placed downstream for capturing the beam leads to a higher transport efficiency. However, the small separation between the source and the solenoid means

that the particles are produced in a region with a non-negligible residual magnetic field strength. Thus, when the particles propagate through the solenoid, the azimuthal velocity component is not fully cancelled by the fringe-fields at the exit plane of the magnet. The beam acquires an angular momentum and rotates and simultaneously expands as it propagates in a field-free region. A nonzero canonical angular momentum gives rise to a centrifugal force and has the same effect as a normalised emittance [127]. For example, a uniform round cold beam of protons emitted along the magnetic field lines has an equivalent magnetisation emittance [249, 250]

$$\epsilon_{\mathrm{rms}} = \frac{eB_s r_s^2}{2\gamma\beta m_p c}, \tag{6.3}$$

where B_s is the residual magnetic field at the source and r_s is the initial radius of the beam. Equation 6.3 is obtained under the assumption of a uniform, zero emittance (no plasma temperature) beam propagating from a round aperture with radius r_s. Furthermore, the emittance contribution from the rotation induced by the magnetic field is expressed in terms of the rms emittance $\epsilon_{\mathrm{rms}} = \sqrt{\langle x'^2\rangle\langle x^2\rangle - \langle xx'\rangle^2}$. The magnetisation plays an important role in the repulsive forces that tend to diverge the beam if the equivalent magnetisation emittance is comparable to the intrinsic emittance arising from the generation process at the source and possible nonlinear effects (e.g. space charge forces).

According to eq. 6.3, a residual magnetic field $B_s = 100\,\mathrm{mT}$ generates an equivalent transverse emittance $\epsilon \sim 10^{-5}\,\mathrm{mm\,mrad}$ for a $10\,\mathrm{MeV}$ proton beam with a radius at the source $r_s \approx 10\,\mu\mathrm{m}$. In this example, the magnetisation emittance is two orders of magnitude lower than typical upper limits measured for proton beams produced by irradiating thin metallic foils with an intense laser [248]. A more general evaluation of the significance of the magnetisation emittance in the case of a proton beam for LhARA is shown in fig. 6.2. Thus, the residual magnetic field at the target is a negligible source of emittance growth compared to the space-charge forces of the protons once they have separated from the co-propagating electron cloud. As a consequence, the distance between the source and a capture solenoid is mainly limited by the space constraints determined, for example, by the transverse size of the solenoid or the layout of the laser-target interaction chamber. A second detrimental effect of a residual field at the particle source is the bending of the target caused by eddy currents possibly induced in a solid tape target. However, this effect was observed to be significant for pulsed solenoids and large residual fields at the target close to $1\,\mathrm{T}$ [251].

6.1.2 Comparison of cable materials

Given the maximum flux density that a solenoid for LhARA has to provide ($\sim 1.3\,\mathrm{T}$), two materials that are conventionally used in similar magnets were considered. The advantages and disadvantages are presented below.

Copper (normal-conducting)

The conventional choice of cable for low-field magnets found in particle accelerators is the hollow copper conductor. The direct contact between the water and the source of heat makes the cooling of the hollow cable

very efficient. Furthermore, the cables are typically potted in resin which leads to a very robust coil. Copper has a higher electrical conductivity than, for example, aluminium, and may lead to overall more compact coils. However, each turn of the coil has to be individually insulated against its neighbouring conductors which requires space and adds complexity in the manufacturing process. Copper also has a higher density which increases the weight of the coil. Lastly, the cooling and electric circuits need to be properly separated at the end of the coil winding resulting in an increased risk of water leakages.

Superconductor

The second conventional technology used for either very high fields or for very large magnets is a superconducting coil. Superconductivity implies zero, or negligible, electrical resistance, so that in DC operation, the magnet does not dissipate power by Joule heating. However, known superconductors work at cryogenic temperatures, so refrigeration power is required to cool down the magnet. As the cryogenic system uses liquid helium, the technology is expensive.

Superconducting coils also suffer from physical limitations. The number of electrons in the superconducting band is not only affected by the temperature of the conductor, but also by the interaction with electromagnetic fields. Rapidly changing magnetic fields can lead to eddy currents and a localised increase in temperature, while strong magnetic fields can change the phase of the conductive material and lead to a reduction in the superconducting electron population. Any too rapid or too strong change in the field of the magnet caused by the interaction of the laser with the target can lead to the magnet quenching, resulting in a large scale Ohmic heating and the loss of superconductivity.

Compared to conventional conductors, superconductors can carry very high current densities, as shown in fig. 6.3, which results in much more compact coils. Although the coil is more compact, a large space is required around it to embed the coil in a rigid coil former and in an insulating cryostat. To prevent the superconductor returning to its resistive state, the cryostat has to be cooled below the critical temperature, T_c, of the conductor material used. Hence, a complex infrastructure is required that has to be maintained carefully. The cost and operation of such an infrastructure leads to an overall budget that is significantly higher than the cost of a conventional magnet.

Superconducting materials are classified in two general categories: Low Temperature Superconductors (LTS), with T_c less than $\sim 35\,\mathrm{K}$, and High Temperature Superconductors (HTS), with T_c above $\sim 35\,\mathrm{K}$. LTS magnets are typically operated at a fixed temperature $\sim 4.2\,\mathrm{K}$, the boiling point of liquid helium. HTS materials open the way to even higher magnetic fields, but they have seen limited applications so far due to their higher cost.

Niobium-titanium, Nb-Ti ($T_c = 9.2\,\mathrm{K}$ [252]), has the lowest critical temperature of the LTS family and it is currently the standard material of the superconducting magnet industry [253] used for example in accelerators or MRI magnets. It is a ductile alloy with mechanical properties which make it easy to turn it into wires and cables ready for winding. Due to its low critical temperature, Nb-Ti requires a more demanding cooling system.
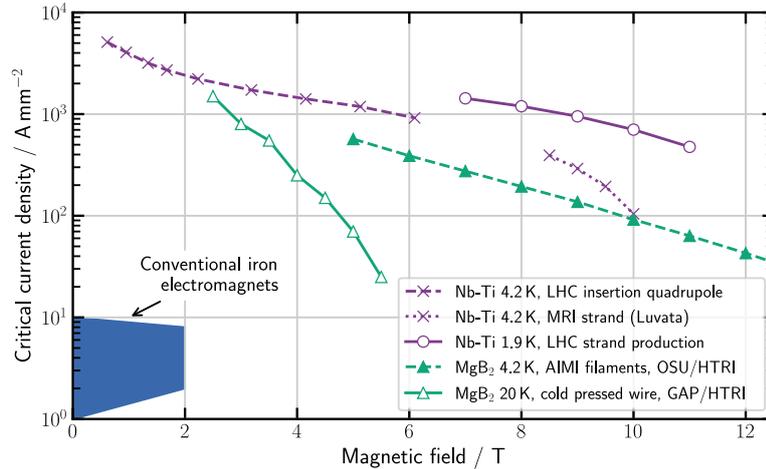
Figure 6.3: The upper limit on the current density (engineering critical current density) that can be carried by a range of superconducting materials as a function of the applied magnetic field. The approximate region of the current densities used in conventional electromagnets is also shown. Experimental data taken from [256] and references therein [257–260].

Higher magnetic fields above $\sim 10\,\text{T}$ can be achieved with the use of the nioboium-tin (Nb_3Sn) compound. The Nb_3Sn wires that are suitable for magnet coils require elaborate assembly and thermal treatment. For this reason, Nb_3Sn wires have only recently transitioned from the R&D phase to series production and are planned for upgrades of large scale accelerator projects such as the High-Luminosity LHC.

Magnesium diboride, MgB_2 ($T_c = 39\,\text{K}$ [252]), is a more recent material and is more suitable for low-field applications in which it can challenge the achievable field of either Nb-Ti or Nb_3Sn at 20–30 K [254]. Its applications have been expanding owing to its lower-cost with respect to HTS materials and its operational margin because of the higher critical temperature and heat capacity. For example, a MgB_2 cable assembly [255] has been produced to deliver a current of about $120\,\text{kA}$ at up to $25\,\text{K}$ to the magnets in the interaction regions of High-Luminosity LHC. Furthermore, the stable operation of a MgB_2 prototype solenoid ($0.3\,\text{m}$, $0.8\,\text{T}$) was demonstrated with an AC plug power $< 3\,\text{kW}$ for cryocooler operation [255].

For LhARA, the choice of material for the solenoids has to be made based on an evaluation of the manufacturing and operational cost, reliability, and ease of operation. The complexities of a superconducting magnet suggest that the main route should be normal-conducting coils. An initial design of a superconducting solenoid for LhARA has been elaborated elsewhere [261]. In the following sections, the aim is to present the design process and results of a normal-conducting solenoid based on hollow copper conductors.

6.1.3   Analytical modelling

An analytical design is necessary to derive the main parameters of a future magnet before starting a more detailed design using numerical methods. An overview of the main calculations done during one iteration of the design in given below. Typically, this set of calculation is repeated until a design is found which achieves the required performance within the typical constraints of magnets that can be operated in practice.

Maximum flux density

The main parameter of a solenoid is the flux density at the centre of the magnet. For a thin-walled solenoid electromagnet of finite length L_s, inner radius r_{coil}, and N turns, the field on the axis is given by

$$B_0 = \frac{\mu_0(NI)}{\sqrt{L_s^2 + 4r_{\text{coil}}^2}}, \tag{6.4}$$

where μ_0 is the vacuum permeability. Thus, the required ampere turns (NI) can roughly be estimated from eq. 6.4 using the length of the solenoid and the required flux density at the centre. The inner radius of the coil is obtained from the aperture matched to the full size of the beam by adding the wall thickness of the vacuum chamber, space for a coil former, ground insulation of the coil and the insulation of the cable.

Although eq. 6.4 provides a formula to estimate the maximum field of a solenoid from the ampere turns (NI), it is only valid for thin coils. The relatively high magnetic fields required by LhARA can be obtained with solenoids that necessarily have a large number of layers and a significant thickness. Thus, a better estimation of the field at the centre of the magnet is obtained by summing the flux density produced by each layer of the coil. Treating a layer as a separate thick solenoid, the flux density produced by the layer at the centre is [262]

$$B_0 = \frac{\mu_0(NI)}{2(r_o - r_i)} \ln \left(\frac{\sqrt{r_o^2 + \left(\frac{s}{2}\right)^2} + r_o}{\sqrt{r_i^2 + \left(\frac{s}{2}\right)^2} + r_i} \right), \tag{6.5}$$

where r_i, r_o, s are the inner radius, outer radius, and length of the coil layer.

Current density

In most normal conducting accelerator magnets, such as dipoles or quadrupoles, the electrical power consumption is proportional to the current density inside the conductor. For solenoids, the outer turns of the coil are significantly longer than the average turn length so that the effect of the reduced current density can be reversed by the disproportionately higher resistance. Once the length of the solenoid and ampere turns (NI) have been fixed, the mass and the thickness of the coil can be evaluated. Figures 6.6 and 6.7 show the expected mass and transverse size of the coil for several practical filling factors. The filling factor f_c is the ratio between the net conductor area and the coil cross section. In the case of hollow conductors, f_c includes the geometric filling factor, the cable insulation, the cooling channel and edge rounding. An optimal structure of the coil has to be found relating to material costs, costs for power supplies. and operating costs.

In order to find the optimal range for the current density j, the total cost of the solenoid was estimated based on general scaling laws adapted from [263]. Taking into account the cost variation from year to year, the costs included in the calculation are normalised to give approximately correct relations between the various types of cost. Thus, the absolute cost remains undetermined. The overall cost was subdivided into investment cost (equipment and magnet capital) and running cost. Figure 6.4 shows that a current density in the range
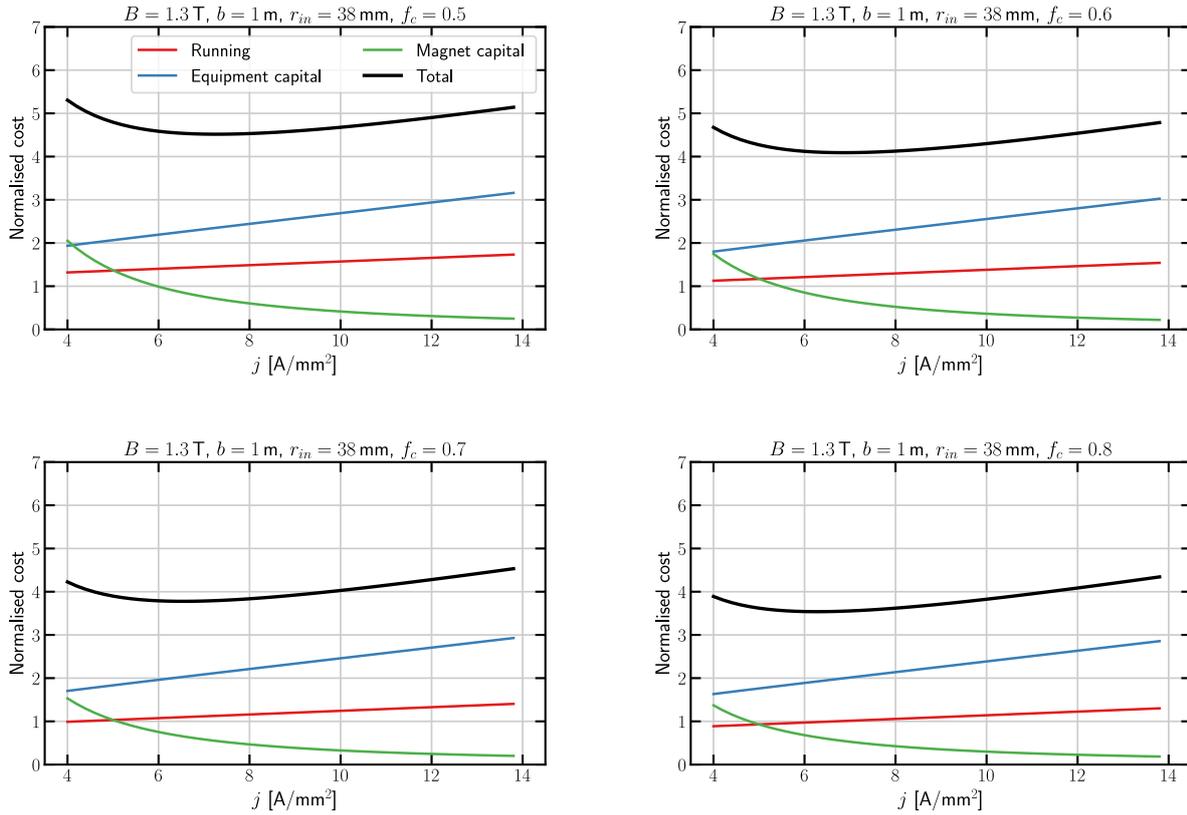
Figure 6.4: The total normalised cost and its components evaluated for a solenoid that is suitable for LhARA as a function of the current density for four different compaction factors f_c. The weight of each component follows a model adapted from [263]. A more compact coil leads to a lower total cost. The current density where the minimum total cost is achieved stays within the same range for all the compaction factors shown here.

6–$8 \, \mathrm{A/mm^2}$ leads to a minimum total cost over the projected magnet lifetime ($60,000$ hours) by balancing the investment (capital) cost against running costs. The current density was restricted to this range during the iterations of the design of the solenoid. As a result, the coils has to be cooled by direct water cooling which is typically employed for current densities $2 \, \mathrm{A/mm^2} < j \leq 10 \, \mathrm{A/mm^2}$. Such a cooling technique requires demineralised water for low electrical conductivity and hollow conductor profiles.

Power consumption

A second critical parameter in the design of the solenoid is the power consumption which drives the running cost and is equal to the power loss in the coil. The total dissipated power per magnet can be estimated from

$$P = \rho j (NI) l_{\mathrm{avg}}, \tag{6.6}$$

where ρ is the resistivity of the coil conductor and l_{avg} is the average turn length. In order to determine l_{avg}, the width of the coil has to be calculated from the coil cross sectional area, A, which, in turn, is fixed by the filling factor f_c according to

$$j = \frac{(NI)}{f_c A} = \frac{I}{a_{\mathrm{cond}}}, \tag{6.7}$$
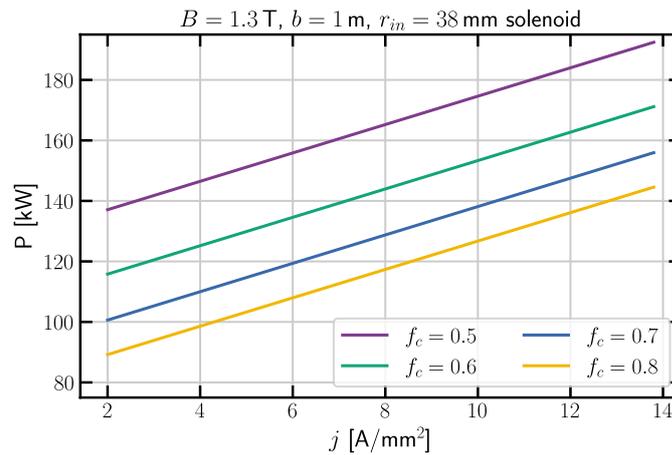
Figure 6.5: Power consumption (at $50\,^{\circ}\mathrm{C}$) of a $1\,\mathrm{m}$ long, $1.3\,\mathrm{T}$ solenoid as a function if the current density j for a range of compaction factors f_c encountered in practice.

where a_{cond} is the conducting cross-section area of the cable. Figure 6.5 shows the scaling of the power consumption of a single solenoid as a function of the current density for several values of the filling factor encountered in practice. Tightly packed coils result in a lower power consumption. However, the size of the channel required for efficient cooling imposes an upper limit on the achievable filling factor.

The resistivity, ρ, of oxygen-free copper changes with temperature and is calculated according to $\Delta\rho = \alpha \cdot \Delta T \cdot \rho(0)$. The base resistivity $\rho(0)$ is taken as $1.715 \times 10^{-8}\,\Omega\mathrm{m}$ [264] (at $T = 20\,^{\circ}\mathrm{C}$) and the temperature coefficient α is $0.004\,29\,^{\circ}\mathrm{C}^{-1}$ [265]. The coil temperature should be kept to a minimum in order to decrease the power consumption. In calculating the cooling requirements below, the temperature of the coil is chosen as $45\,^{\circ}\mathrm{C}$ to obtain upper limits for the power consumption and the resistance of each respective design.

Leading up to a full set of magnet specifications, it is useful to inspect how the main parameters of interest of the coil scale within the range of typical current densities for water-cooled copper conductors. Figure 6.5 shows that a power consumption in excess of $100\,\mathrm{kW}$ is expected. The changes in the total mass and the thickness of the coil with the current densities are shown in figs. 6.6 and 6.7, respectively. For a current density around $8\,\mathrm{A/mm^2}$, the coil is expected to weigh about one tonne and have an outer diameter between $20\,\mathrm{cm}$ and $30\,\mathrm{cm}$.

Cooling requirements

In the case of water-cooled magnets, the heat is removed by water circulating in the hollow conductor of the coil. The choice of cooling parameters and number of cooling circuits is based on several principles. A complete list of the practical recommendations and conventional values [245, 246] applied in the design of the solenoid is included in table 6.1.

For example, an excessive water velocity has to be avoided to limit the erosion and corrosion of the cooling pipes and the onset of vibrations. Furthermore, the water flow has to be moderately turbulent for higher thermal efficiency and the flow rate has to lie within the limits of realistic cooling plants. A maximum permitted temperature of less than $60\,^{\circ}\mathrm{C}$ on the coil surfaces is considered to be good practice [246] as it reduces the
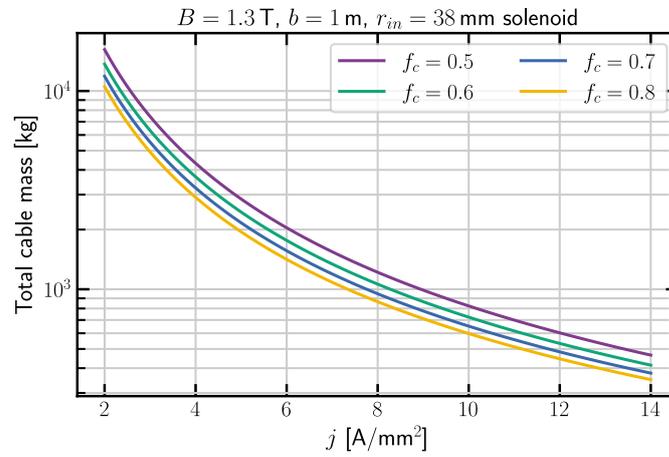
Figure 6.6: The mass of a $1\,\mathrm{m}$ long coil that produces a flux density of $1.3\,\mathrm{T}$ as a function of the current density j for several compaction factors f_c encountered in practice.
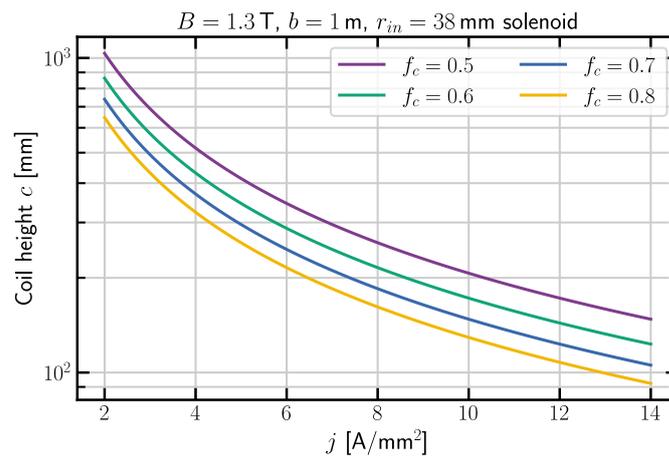
Figure 6.7: The thickness of a $1\,\mathrm{m}$ long, $1.3\,\mathrm{T}$ solenoid as a function of the current density j and the compaction factor f_c.

accelerated ageing of the insulation materials especially in the presence of ionising radiation.

The quality of the water is also essential for the reliability of the coil and for a low rate of ageing due to erosion and corrosion. The resistivity of water needs to be maintained $>0.1 \times 10^6 \, \Omega\text{m}$, its pH between 6 and 6.5, and the level of dissolved oxygen below $0.1 \, \text{ppm}$. Filters are also necessary to remove particles or loose deposits that could obstruct the cooling channels.

In the following, the formulas required to calculate the parameters of the cooling circuit are presented. The cooling pipes are assumed to be long, straight, and smooth with the temperature of the inner surface of the conductor equal to that of the coolant. Furthermore, the conductor is assumed to be isothermal over its entire cross-section.

Table 6.1: Conventional values used in the design of water cooling circuits for normal-conducting magnets as recommended in [245] and [246].

Parameter	Range
Pressure drop	$1 \, \text{bar} \leq \Delta p \leq 10 \, \text{bar}$
Temperature rise	$\Delta T \leq 30 \, °\text{C} \quad (\Delta T \leq 15 \, °\text{C})^\dagger$
Water velocity	$v < 3 \, \text{m/s}$
Reynolds number	$4000 < Re < 10^5$

† for advanced stability

In practice, each layer of the coil corresponds to a separate cooling circuit. Thus, the outmost layer of the coil is also the longest cooling circuit and presents the most demanding cooling requirements. As the length of the cable in a layer increases with the square of the radius of the turns in the layer, the cooling circuits closest to the bore of the magnet and those furthest away from the axis have significantly different lengths and amounts of power that need to be transferred. If several circuits are connected in parallel, the pressure drop will be the same across each of the circuits and the flow will be shared according to the hydraulic resistance of each circuit.

Considering one layer of the coil, the power transfer P in the equivalent cooling circuit determines the required fluid flow rate which in practical units is [245]

$$Q \left[\frac{\text{L}}{\text{min}} \right] = 14.3 \frac{P[\text{kW}]}{\Delta T}, \tag{6.8}$$

where ΔT is the temperature increment along the cooling pipe. Then, the pressure drop Δp can be estimated from Blasius law by assuming a hydraulically smooth pipe of length l_{cc} and diameter d from [245]

$$\Delta p = 60 \, l_{\text{cc}}[\text{m}] \frac{Q \left[\frac{\text{L}}{\text{min}} \right]^{1.75}}{d[\text{mm}]^{4.75}}. \tag{6.9}$$

If the assumption of smooth pipes is lifted, a dimensionless friction factor f has to be introduced in the right hand side of eq 6.9. For turbulent flow, the friction factor is transcendental [246] and it is defined by the Prandtl-

Colebrook expression

$$\frac{1}{\sqrt{f}} = -2\log_{10}\left(\frac{\epsilon}{2.71d} + \frac{2.51}{Re\sqrt{f}}\right) \tag{6.10}$$

where ϵ is a measure of the roughness of the cooling channel (typically 1.5×10^{-3} mm), d is the inner pipe diameter and Re is the Reynolds number. The hydraulics calculations can be done with dedicated software tools that are based on eq. 6.10. In the design iterations described here, such a tool [266] was used to confirm that the simplified treatment of eq. 6.9 does not lead to large deviations. It is also important to note that the pressure drop is approximately proportional to $1/d^5$. Thus, increasing the cooling channel diameter by a small factor can reduce the required pressure drop significantly.

In designing a solenoid, a more useful formula can be obtained by combining eqs. 6.8 and 6.9. The power, P, that needs to be removed dictates the required diameter of the cooling channel as [246]

$$d = 5 \times 10^{-4} \left(\frac{P}{\Delta T}\right)^{0.368} \left(\frac{l_{\text{cc}}}{\Delta p}\right)^{0.21}, \tag{6.11}$$

It is important to note that the power, P, dissipated in the cable depends on the length of the cable l_{cc}, which in the process of designing a solenoid is influenced by the overall transverse size of the cable, and thus, by the diameter d of the cooling channel. If a parametrization $l_{\text{cc}} = l_{\text{cc}}(d)$ is found, eq. 6.11 can be treated as a nonlinear relation that allows one to find the diameter of a cooling channel given the maximum allowed ΔT and Δp. Lastly, the Reynolds number, representing the ratio between the inertial forces and the viscous forces, can be calculated from

$$Re = \frac{dv}{\nu}, \tag{6.12}$$

where v is the water speed and ν is the dynamic viscosity of water. A high Reynolds number ($Re > 4000$) indicates turbulent flow and a more efficient collection of heat from the inner surface of the pipe compared to laminar flow.

For the solenoid required by LhARA, a general conclusion is that for a fixed current density and filling factor, a high current is necessary to increase the total cross-section of the cable and allow for enough space for the cooling channel. The required minimum current can be lowered by the use of more cooling circuits. For a low filling factor and practical pressure drop and temperature increase, the minimum current required is at least a few hundred Amperes.

Separation of electric and coolant circuits

In a hollow-copper conductor, both the current and the cooling fluid flow through the same cable. If multiple layers of the coil are managed by the same cooling circuit, the fluid has to be fed in parallel into the layers, while the electrical connection has to be in series. Thus, the large number of layers leads to an intricate piping scheme. The solution is to use distributing manifolds.

Figure 6.8 shows a manifold that consists of an electrically non-conductive body that ensures the inflow
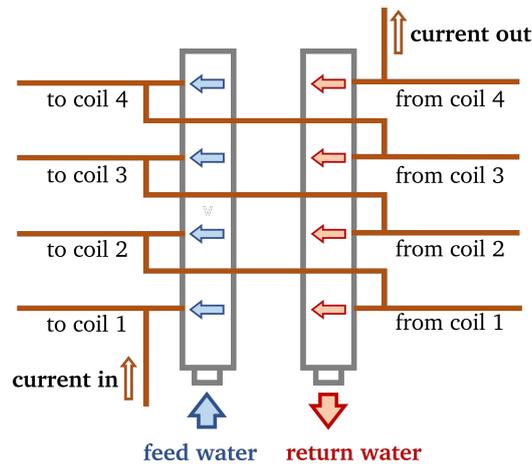
Figure 6.8: Schematic representation of two cooling manifolds (grey colour) that separate the coolant flow from the electric connections. Water is fed in parallel through the first manifold to four layers of the solenoid. A second manifold ensures the return of the coolant. The flow of current is also indicated; adapted from [267].

(or outflow) of the cooling fluid [267]. The individual ends of the coil layers are connected to one manifold receiving the fluid in parallel. A second identical manifold ensure the return flow of the coolant. Thermoplastic polymers such as PEEK (polyether ether ketone) are used as the material of the manifolds because they are water-resistant over time and electrically isolate the coil layers. An electric connection is added between the two manifolds from one layer of the coil to the next one.

Mechanical resistance

In view of the typical next stage in the design of the solenoid for LhARA, the mechanical design, special consideration has to be given to the structural resistance of the coil. Additional reinforcement of the coil can be obtained with a tube through the bore of the solenoid or an outer tube to strengthen the structure. The magnetic pressure inside a cylindrical magnet is defined by the flux density B at the centre as [268]

$$P_{\text{mag}} = \frac{B^2}{2\mu_0} \tag{6.13}$$

For a $1.3\,\text{T}$ solenoid, the magnetic pressure is approximately $P_{\text{mag}} \approx 670\,\text{kPa}$.

The internal pressure creates a hoop stress on the pipe. The hoop stress σ_θ is the stress around the circumference of a cylinder and it is caused by a radial pressure gradient. For a thin-walled ($r_m/t > 10$) pressure vessel, such as one layer of the solenoid, Barlow's formula is typically employed to calculate the hoop stress as [269]

$$\sigma_\theta = \frac{r_m}{t} P_{\text{mag}}. \tag{6.14}$$

For the solenoids required by LhARA, the copper windings have a radial thickness $t \approx 1\,\text{cm}$ and a mean radius $r_m \approx 24\,\text{cm}$. For these average coil dimensions, the solenoids for LhARA should withstand a stress $\sigma_\theta \approx 16\,\text{MPa}$. To determine the mechanical load on the solenoid windings, the hoop stress has to be compared to the yield strength (35–65 MPa for soft oxygen-free copper, Cu-OF [264]) and the ultimate tensile strength (200–
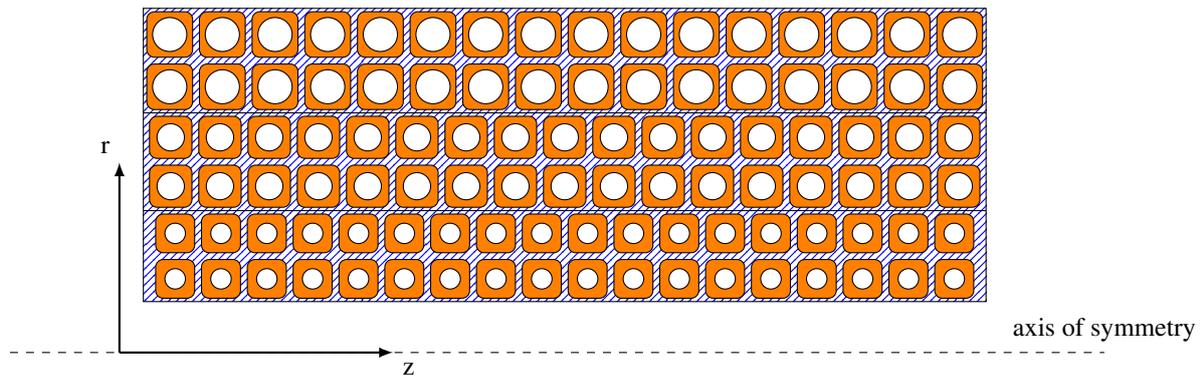
Figure 6.9: Schematic representation of a graded coil design. The transverse cut through the coil shows three different types of hollow copper cables in the upper half of the solenoid. The cable with the smallest transverse size is used for the first two layers counting from the bore of the magnet. The diameter of the cooling channel increases toward the exterior of the coil to balance the larger power loss and pressure drop in the respective layers.

220 MPa for Cu-OF [264]). Therefore, based on this simplified calculation, the coil will not experience plastic deformations or structural damage. However, a more elaborate mechanical modelling beyond the analytical treatment is required if a design of the solenoid will be further developed for LhARA. Typically, the mechanical deformations of the coil are simulated in 2D or 3D using the finite element method (FEM) and a dedicated software that can couple the generation of forces and stress to the heat and current flows through the solenoid.

The basic principles of normal-conducting magnet design have been presented here to give the reader an understanding of the large number of parameters that have to be included in an analytical model. In addition, the conventional values used in practice have been quoted. These practical constraints were combined with the flux density and length required by a solenoid for LhARA to evaluate the performance of the magnet over the range of current densities typical for water-cooled copper coils.

6.1.4   Design results

As a result of the large typical coil thickness found in the design iterations for a solenoid, the overall transverse size of the cable and the cooling pipe are dictated by the proper heat removal from the outmost layer of the coil. This approach leads to an overdesign of the cooling of the cables closer to the bore of the magnet since they have a reduced length and power dissipation. A superior coil design can be obtained by grading—using wires with a smaller cross-section closer to the axis of the coil with the same conducting area, but a narrower cooling pipe. In this way, more turns can be fitted in the layers closer to the bore and a higher flux density can be achieved at the centre of the magnet. Using three different cable specifications was found to be optimal, as a larger number leads to a higher manufacturing complexity for the cables with minor changes in the performance of the magnet. A schematic of the structure of a graded coil design is shown in fig. 6.9.

A methodology was devised to calculate the parameters of interest of a particular solenoid design by executing the following steps. The solenoid is built from three different types of hollow copper cables. In the three sets of layers corresponding to the three types of cable, the optimal transverse size of the cable is found such that the
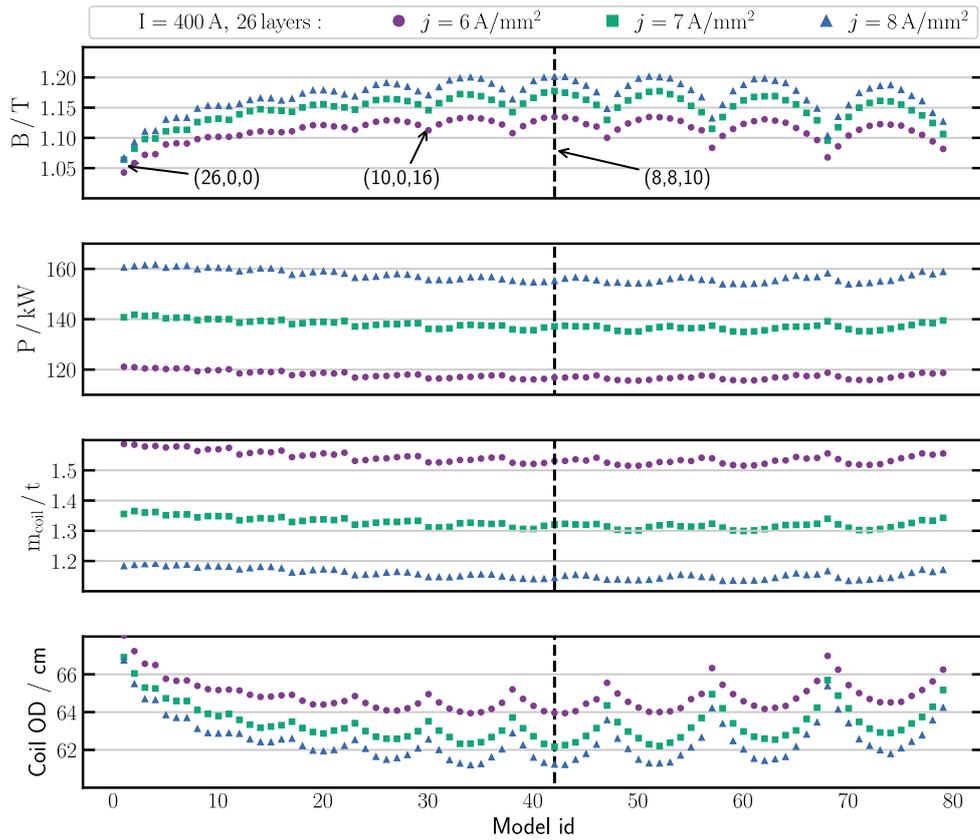
Figure 6.10: The main performance metrics of all the possible configurations of a 400 A graded solenoid with 26 layers and three different types of copper cable. Each point represents a different configuration. The numbers in the round brackets indicate the number of layers made from each cable type starting from the bore of the magnet. The dashed line marks the analytical model that produces the highest flux density.

outmost layer in that set can be cooled efficiently. Here, efficient heat removal is considered to correspond to a temperature increment $\Delta T = 30\,°C$ and a pressure drop around a maximum $\Delta p = 8\,\text{bar}$. Equation 6.11 is used to find the optimal diameter of the cooling channel with $0.1\,\text{mm}$ precision. Given these constraints, an analytical model of the magnet can be built starting from the bore and adding layers of cable given the current density, total current, length, and total number of layers of the solenoid. Only even number of layers were considered in each group to allow adjacent layers to be wound in opposite directions such that any transverse field component is minimised.

For a fixed number of total layers, the calculations described above are carried out for all possible coil configurations. Then, the maximum flux density, power, size, and mass of the configurations are compared and a final design is chosen. Three nominal currents were chosen based on power supplies that are known to be in use in accelerator facilities.

400 A solenoid

Figure 6.10 shows the key parameters of a $1\,\mathrm{m}$ long $400\,\mathrm{A}$ solenoid with 26 layers according to analytical calculations. The trends in fig. 6.10 show that an optimal configuration exists which maximises the flux density at the centre of the magnet. The gain in the flux density is a result of a smaller transverse size of the coil which overall brings the layers closer to the central axis. However, the minima of the power consumption or the total mass of the coil do not correspond to the same coil configuration that leads to the maximum flux density. Since the relative gain in flux density is larger compared to the relative decrease in the power consumption, the final design is chosen by favouring the maximum achievable magnetic field. A more complete list of specifications is provided in table 6.2.

Figure 6.10 also shows that a graded coil design reduces the length of conductor required by improving the maximum achievable flux density and decreasing the power consumption. The change from one type of cable to three types results in a 13% increase in the flux density at the centre and a drop of 4% in both the power consumption and mass of the coil. It is important to note that a higher current density leads to a small relative increment in the maximum flux density compared to the increase in power consumption. The route to higher fields at constant current, while balancing the power loss, is to use more layers. However, the complexity of the cooling system rises with the number of layers and the magnet becomes less reliable.

1000 A solenoid

All the allowed coil configurations were also calculated for a $1\,\mathrm{m}$, $1000\,\mathrm{A}$ solenoid with 14 layers and the results are shown in fig 6.11. Since the total number of layers is lower than that used for the $400\,\mathrm{A}$ coil, the benefit of using a graded configuration is smaller. At $8\,\mathrm{A/mm^2}$, the maximum flux density can be increased by 6% while the power consumption and the mass can be lowered by less than 1%. However, the cooling requirements are still better balanced between the different layers of the coil if more than one type of cable is used. A more detailed list of specifications for the coil that generates the maximum field at the centre is included in table 6.3. The specifications of a design with 10 layers are presented in table 6.4.

2000 A solenoid

A high current, $2000\,\mathrm{A}$, solenoid was also considered. At a higher current, fewer turns are needed for a fixed value of the magnetic field at the centre. However, if the current density is not increased in proportion to the current, the transverse cross-section of the cable needs to be enlarged as well. This leads to a similar overall size of the coil comparable to the $1000\,\mathrm{A}$ design, as indicated by fig 6.12. Here, the optimal design was chosen to have 10 layers and the current density was kept in the same range as for the previous iterations to keep the power consumption down at $\sim 105\,\mathrm{kW}$. A more compact version can be achieved with 8 layers and a current density $\sim 11\,\mathrm{A/mm^2}$ that leads to a power consumption close to $\sim 130\,\mathrm{kW}$.
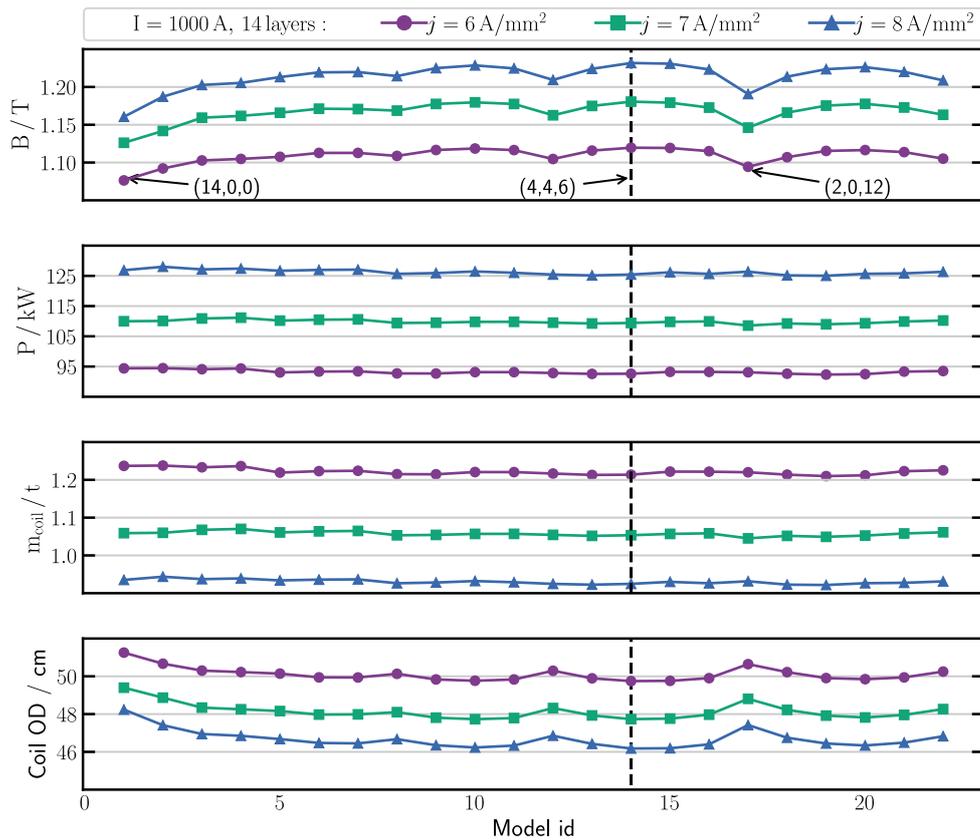
Figure 6.11: The main performance metrics of all the possible configurations of a 1000 A graded solenoid with 14 layers and three different types of copper cable. Each point represents a different configuration. The numbers in the round brackets indicate the number of layers made from each cable type starting from the bore of the magnet. The dashed line marks the analytical model that produces the highest flux density. A line is added between the markers to guide the eye.
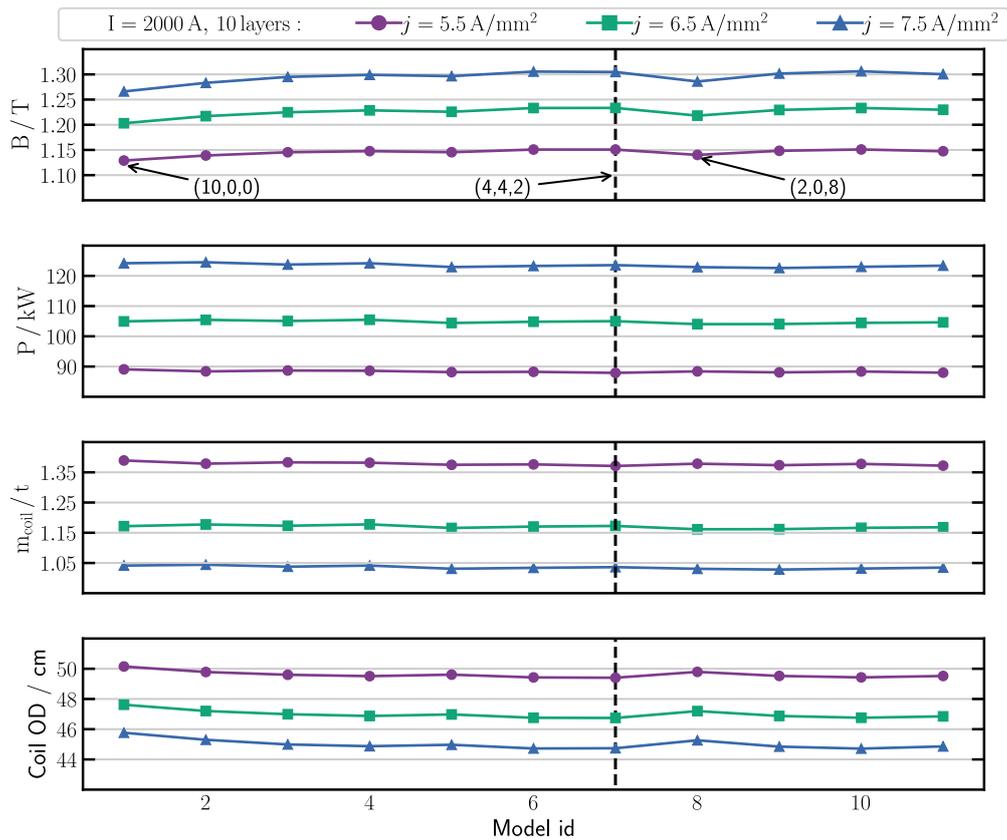
Figure 6.12: The main performance metrics of all the possible configurations of a 2000 A graded solenoid with 10 layers and three different types of copper cable. Each point represents a different configuration. The numbers in the round brackets indicate the number of layers made from each cable type starting from the bore of the magnet. The dashed line marks the analytical model that produces the highest flux density. A line is added between the markers to guide the eye.

6.1.5 Numerical 2-D simulations

The typical design stage that follows the analytical calculations for a magnet is to compute its performance with numerical simulations. The program used to calculate the maximum achievable flux density, the integrated magnetic field, and the stored energy is described below together with a summary of the numerical results.

The designs specified in tables 6.2, 6.3 and 6.4 were modelled in 2-D with the use of FEMM [223]. Since the solenoid has axial symmetry, 2-D simulations are faster and provide a good first set of results before full 3-D simulations are done. The program suite FEMM can solve nonlinear magnetostatic problems and provides a post-processor to evaluate different quantities of interest at a specified point, along a contour, or inside a volume. The problem is defined from the geometry of the elements by defining the material properties, the boundary conditions, and the current distributions.

A magnetostatic problem is a problem in which the fields are time-invariant. In this case, the current density, J, acts as the source of the field intensity, H, and the flux density, B, which must obey

$$\nabla \times H = J \tag{6.15}$$

$$\nabla \cdot B = 0 \tag{6.16}$$

where B and H are related through $B = \mu H$. The magnetic permeability μ is a characteristic of each material. If a material is nonlinear (e.g. saturating iron), the permeability, μ, is actually a function of B:

$$\mu = \frac{B}{H(B)}. \tag{6.17}$$

A flux density that satisfies eq. 6.16 can be written in terms of a magnetic vector potential, A, as $B = \nabla \times A$. Then, eqs. 6.15 and 6.16 are satisfied by a solution of the equation

$$\nabla \times \left(\frac{1}{\mu(B)} \nabla \times A \right) = J. \tag{6.18}$$

Thus, FEMM solves a single equation, 6.18, using a triangular mesh over the simulation region and then computes B and H from A. In a 2-D axisymmetric problem, the vector A has only one component in the "out of the page" direction.

In the numerical simulations described here, the material of the coil was taken as linear copper (i.e. relative permeability $\mu_r = 1$). The shield around the coil is made out of nonlinear pure annealed iron. The $B-H$ curve used in the models is the one of the basic data provided by FEMM and showed in fig. 6.13. To build some tolerance into the model, the thickness of the iron shield was chosen such that the maximum field strength in the iron is around 1.6 T, which is closer to the saturation value for cast iron.

The results of the magnetic field calculations are summarised in fig. 6.14 which shows the magnitude of the flux density for the three designs of a solenoid corresponding to three different currents. In each case, the
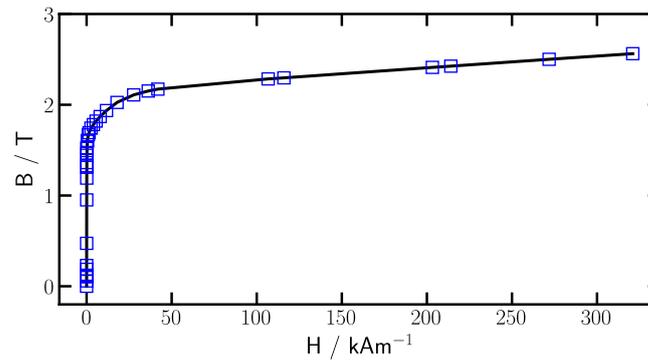
Figure 6.13: Magnetic flux density as a function of the field intensity showing the nonlinear permeability of pure iron, annealed. The data is provided as part of FEMM [223].

simulations include an iron flux return. The field distributions are used to extract the flux density at the centre, the integrated field strength, and the stored energy for each design.

Further optimisation of the shape of the flux return and the ends of the coil or the spacing between them is required to ensure that enough space is available to bring the cables from each layer outside of the coil into the manifold that provides the water input and output. This is typically part of the modifications to the main design that are done by the coil manufacturer and is out of the scope of the work presented here.

A comparison between the magnetic fields on the axis of the solenoids for the three nominal currents is shown in fig. 6.15. The flatness of the field inside the solenoids is comparable for the three currents. Furthermore, due to the iron cover, the field drops relatively quickly with distance from the edge of the magnet. For example, at $10\,\mathrm{cm}$ from the entrance of the magnet, the flux density is below $10\,\mathrm{mT}$.

6.2   Discussion

A **preliminary design for a normal-conducting solenoid** was proposed in this chapter and analytical models were presented that evaluate the parameters of the magnet and its capability to produce the strong focusing field required for LhARA. The design uses **hollow copper cables and water cooling** which can facilitate the production of such a solenoid magnet. The normal-conducting (NC) technology was chosen as the first case for a design study since NC magnets are easier to be integrated and run in a facility due to the absence of the cryogenic system encountered in superconducting solenoids. As LhARA plans to deliver proton and ion bunches at a repetition rate of $10\,\mathrm{Hz}$, the magnet is intended to be used in continuous-wave mode and not to be pulsed.

The key feature of the main design is the **more compact arrangement of the layers of cable obtained by grading**–using cables with different cross-section, but the same conducting area, in different sets of layers of the coil. Compared to the case of using a single type of cable, the graded coil design results in **smaller size and total mass of the magnet, and in more uniform cooling** requirements across the distinct layers of the coil.

The main solenoid design was adapted to three cases, each corresponding to a different current source that could be readily available or preferred for LhARA, producing $400\,\mathrm{A}$, $1000\,\mathrm{A}$, or $2000\,\mathrm{A}$. The two largest
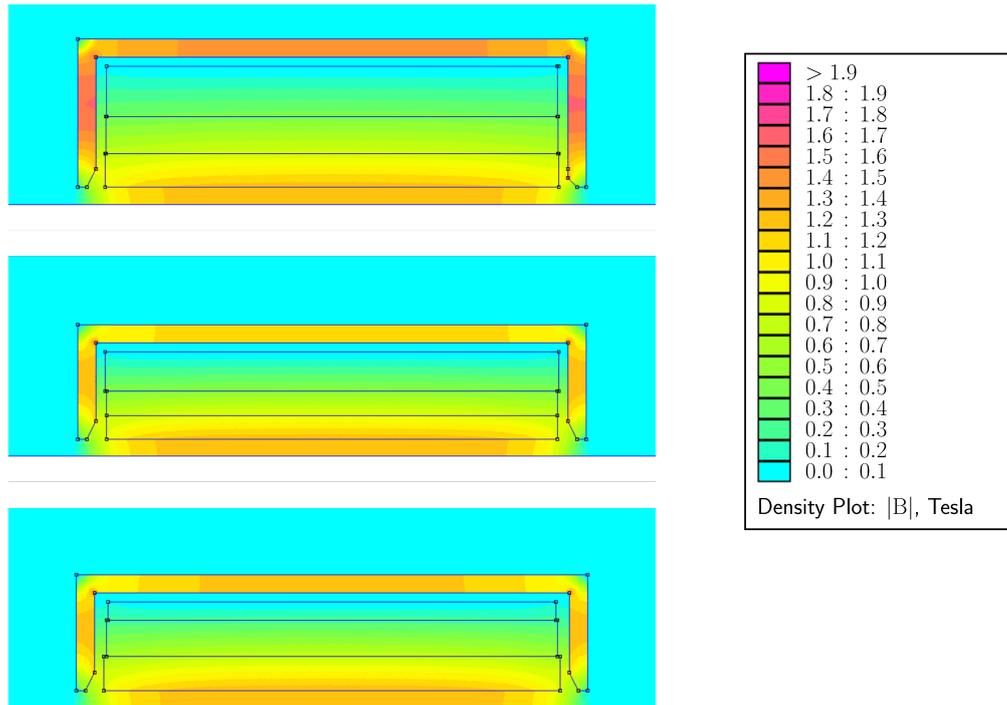
Figure 6.14: The magnitude of the flux density shown in a longitudinal cross-section through the three design versions of the solenoid for LhARA: 400 A (top), 1000 A (middle), 2000 A (bottom). Both the contour of the coil and of the iron shield around the coil are indicated. Each coil is composed of three different regions. The transverse size of the copper cable differs between the separate regions of the same coil.
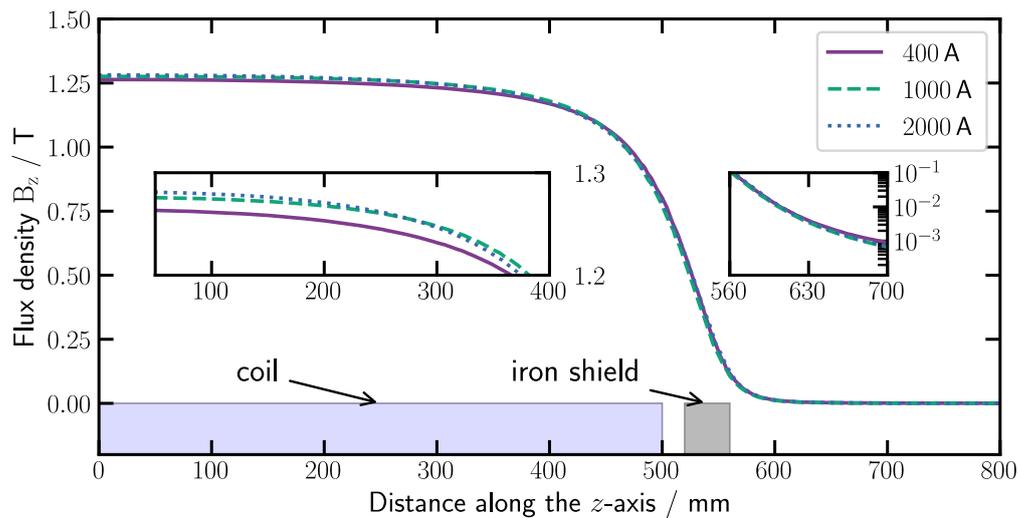
Figure 6.15: The longitudinal component of the flux density on the axis of the solenoid obtained from 2D magnetostatic simulations with FEMM [223]. The value of the current in the legend indicates the corresponding design of the solenoid. The inset on the left shows the flatness of the field inside the magnet. The inset on the right shows the magnitude of the flux density immediately outside of the iron shield.

currents result in the designs with smaller total mass and power consumption. The length of the coil is 1 m in all the designs. The total **mass of the magnet is in the range 1500–**2000 **kg**, while the **power consumption varies between 105** kW **and 155** kW, and goes down for higher nominal current.

One factor that was identified to influence the reliable use of the solenoids as part of a research facility is the **intricate piping scheme** which is **needed to ensure the inflow and outflow of the cooling liquid**. Two distributing manifolds made of thermoplastic polymers were proposed in addition to electrical connections between the layers of the coil. Additional space is required at the two ends of the magnet and on its side to bring the cooling pipes out from the coil and guide them to the manifolds. The magnet design based on a current of 2000 A contains the smallest number of layers of cable and, thus, would require the least complex piping scheme.

Another aspect relevant in this context is the residual magnetic field produced by the solenoid at the location of the tape target due to the short distance ($\sim$10 cm) between the source and first solenoid required to capture efficiently the protons and ions accelerated from the target. The contribution of the magnetic field to the total emittance of the beam at the source was briefly explained in section 6.1.1. The emittance growth induced by the residual magnetic field was evaluated to be negligible compared to the transverse emittance with which the protons are generated at the source. The close proximity of a solenoid to a tape target has been reported [251] to cause bending of the tape and an increase in the divergence angle of the accelerated protons for residual magnetic fields of the order of $\sim$1 T. In the designs presented in this chapter, the iron shield placed around the coil of the magnet reduces the **residual flux density to below 10** mT **at approximately 10** cm **away from the entrance of the magnet**. Thus, in principle, the first solenoid in LhARA could be placed closer to the target to increase the capture efficiency. However, other constraints, such as the layout of the laser-target interaction chamber, transverse size of the solenoid, or the space required for diagnostics of the laser or the target must be considered to decide whether the first solenoid can be placed closer to the target compared to current separation defined in the baseline design of LhARA [270].

As presented in section 6.1.3, a **series of analytical models were used to generate the designs for the solenoids** that could produce the required field strength and be effectively cooled down in practice. The models were **based on practical recommendations and conventional values** [245, 246] used in the planning and construction of magnets for particle accelerators. Moreover, **2-D magnetostatic simulations were carried out to validate** that the solenoids can produce the required magnetic field strength and quality. In order to improve and further validate the designs presented here, 3-D simulations of the current flow, magnetic field, and heat conduction are required.

In conclusion, this chapter presented a **preliminary solenoid design capable of producing the focusing strength required to capture and transport protons and ions from a laser-driven source in LhARA**. It was shown that a suitable magnet could be built based on **conventional normal-conducting technology** that

would result in **moderate to high overall size, mass, and power-consumption**. Even though some of these numbers are typical for magnets used at large particle accelerator facilities, several of such solenoids may not provide the optimal balance between the overall size, power consumption and their capability to generate strong focusing and efficient beam capture, especially in the view of developing a future clinical facility based on the laser-hybrid approach. Furthermore, the extension of the laser-hybrid technique to deliver proton and ion beams for clinical use would require high-field solenoids that cannot be based on normal-conducting technology.

Solenoid magnets are also considered in the next chapter which presents a study comparing the capabilities to capture and transport protons, as well as to reduce the initial energy spread of the beam, using solenoids and plasma lenses in Stage 1 of LhARA.

Table 6.2: Design specifications of a 400 A solenoid for LhARA.

Parameter	Value			Units
Flux density at the centre[†]			1.264	T
Field integral along $(0, z)$[†]			1.262	Tm
Nominal current			400	A
Total number of turns			2514	
Number of layers			26	
Current density			8	A/mm^2
Vacuum bore radius			30	mm
Coil inner radius			38	mm
Coil thickness			268	mm
Coil length			1.00	m
Layer number[‡] (counting from the bore)	1–8	9–16	17–26	
Number of turns per layer	106	97	89	mm
Hollow-conductor size (copper only)	8.4×8.4	9.2×9.2	10.1×10.1	mm^2
Diameter of the cooling channel	5.0	6.6	8.1	mm
Total water flow rate	1.6	3.4	4.9	L/min
Temperature rise	30	30	30	°C
Pressure drop*	7.8	7.6	8.0	bar
Pressure drop**	4.0	6.5	6.8	bar
Water velocity	1.3	1.7	1.6	m/s
Reynolds number	11000	18000	21000	
Cable length			2555	m
Total copper mass			1145	kg
Iron cover mass			850	kg
Electric resistance (at 45 °C)			970	mΩ
Inductance			369	mH
Ohmic power loss (at 45 °C)			155	kW
Total stored energy[†]			29.5	kJ
Width of conductor insulation			0.5	mm
Coil ground insulation width			2	mm
Cable edge radius			1	mm
Iron cover thickness			40	mm
Gap between coil and yoke (r and z)			20	mm
Chamfer width radial, longitudinal (r, z)			$(40, 20)$	mm
Winding bed thickness			4	mm

[†]The values were computed in a 2D finite-element magnetostatic simulation of the magnet.
[‡]Each layer is cooled by a separate cooling circuit. The cooling parameters correspond to the outmost layer in the indicated section of the coil, which is the layer with the highest cooling requirements.
*Blasius law (eq. 6.9), hydraulically smooth pipe
**pressure-drop calculator [266], pipe roughness 0.0015 mm, water at 45 °C

Table 6.3: Design specifications of a 1000 A solenoid for LhARA.

Parameter	Value			Units
Flux density at the centre[†]			1.276	T
Field integral along $(0, z)$[†]			1.1269	Tm
Nominal current			1000	A
Total number of turns			1010	
Number of layers			14	
Current density			8	A/mm^2
Vacuum bore radius			30	mm
Coil inner radius			38	mm
Coil thickness			193	mm
Coil length			1.00	m
Layer number[‡] (counting from the bore)	1–4	5–8	9–14	
Number of turns per layer	76	73	69	mm
Hollow-conductor size (copper only)	12.0 × 12.0	12.6 × 12.6	13.4 × 13.4	mm^2
Diameter of the cooling channel	4.9	6.4	8.3	mm
Total water flow rate	2.2	4.6	7.1	L/min
Temperature rise	30	30	30	°C
Pressure drop*	8.3	8.3	7.8	bar
Pressure drop**	4.3	7.1	6.7	bar
Water velocity	2.0	2.4	2.2	m/s
Reynolds number	16000	26000	30000	
Cable length			826	m
Total copper mass			925	kg
Iron cover mass			645	kg
Electric resistance (at 45 °C)			125	mΩ
Inductance			38	mH
Ohmic power loss (at 45 °C)			125	kW
Total stored energy[†]			19.0	kJ
Width of conductor insulation			0.5	mm
Coil ground insulation width			2	mm
Cable edge radius			1	mm
Iron cover thickness			40	mm
Gap between coil and yoke (r and z)			20	mm
Chamfer width radial, longitudinal (r, z)			(40, 20)	mm
Winding bed thickness			4	mm

[†]The values were computed in a 2D finite-element magnetostatic simulation of the magnet.
[‡]Each layer is cooled by a separate cooling circuit. The cooling parameters correspond to the outmost layer in the indicated section of the coil, which is the layer with the highest cooling requirements.
*Blasius law (eq. 6.9), hydraulically smooth pipe
**pressure-drop calculator [266], pipe roughness 0.0015 mm, water at 45 °C

Table 6.4: Design specifications of a 2000 A solenoid for LhARA.

Parameter	Value			Units
Flux density at the centre[†]			1.281	T
Field integral along $(0, z)$[†]			1.274	Tm
Nominal current			2000	A
Total number of turns			506	
Number of layers			10	
Current density			6.5	A/mm^2
Vacuum bore radius			30	mm
Coil inner radius			38	mm
Coil thickness			196	mm
Coil length			1.00	m
Layer number[‡] (counting from the bore)	1–4	5–8	9,10	
Number of turns per layer	52	50	49	mm
Hollow-conductor size (copper only)	18.2×18.2	18.7×18.7	19.0×19.0	mm^2
Diameter of the cooling channel	5.4	7.3	8.2	mm
Total water flow rate	3.0	6.9	8.2	L/min
Temperature rise	30	30	30	°C
Pressure drop*	8.2	8.2	7.6	bar
Pressure drop**	4.0	7.0	6.6	bar
Water velocity	2.2	2.8	2.6	m/s
Reynolds number	20000	33000	35000	
Cable length			425	m
Total copper mass			1173	kg
Iron cover mass			655	kg
Electric resistance (at 45 °C)			26	mΩ
Inductance			10	mH
Ohmic power loss (at 45 °C)			105	kW
Total stored energy[†]			20.7	kJ
Width of conductor insulation			0.5	mm
Coil ground insulation width			2	mm
Cable edge radius			1	mm
Iron cover thickness			40	mm
Gap between coil and yoke (r and z)			20	mm
Chamfer width radial, longitudinal (r, z)			$(40, 20)$	mm
Winding bed thickness			4	mm

[†]The values were computed in a 2D finite-element magnetostatic simulation of the magnet.
[‡]Each layer is cooled by a separate cooling circuit. The cooling parameters correspond to the outmost layer in the indicated section of the coil, which is the layer with the highest cooling requirements.
*Blasius law (eq. 6.9), hydraulically smooth pipe
**pressure-drop calculator [266], pipe roughness 0.0015 mm, water at 45 °C

Chapter 7

Performance evaluation of the plasma lenses for LhARA

Another subject addressed in this doctoral thesis is the performance of the electron plasma lenses as focusing elements of the LhARA beam-lines for laser-driven proton beams. As discussed in section 1.1, protons generated via the TNSA mechanism exhibit particular features such as wide energy spread, large divergence at the source, and small transverse emittance and initial beam size. Each of these properties impact the evolution of the beam along the LhARA beam-lines and has to be considered in the development of LhARA towards the delivery of highly uniform, parallel beams of variable transverse size at the *in-vitro* end station. Moreover, the beam-line must allow protons within the desired energy range of $15\,\mathrm{MeV}\pm2\%$ to be selected from the wide energy spectrum generated at the source. As part of the energy selection scheme, it is desirable for most of the protons within the nominal energy range produced at the source to be transported to the end station with minimal beam losses.

The electron plasma lenses have favourable properties that allow, in principle, efficient capture, transport, and energy selection of protons from a laser-driven source. For example, the energy-dependent focal length of the plasma lenses allow for protons to be filtered based on their kinetic energy with the use of a beam collimator. Furthermore, a plasma lens can be designed with a larger transverse aperture without a significant increase in the overall size or cost, as would be the case for a normal-conducting (NC) solenoid. A design for a NC solenoid as an alternative to the plasma lenses for LhARA is presented in chapter 6. Such a conventional magnet was shown to provide equivalent focusing strengths as the plasma lenses envisioned for LhARA, though at the expense of larger size and cost. A comparison of the beam-focusing capabilities of a solenoid and a plasma lens has been presented in section 1.3.4 based on the analytic expressions of their focusing strengths. For the results presented in this chapter, comparative simulations of the LhARA Stage 1 beam-line with plasma lenses and solenoids were performed and analysed.

In order to evaluate the performance of the electron plasma lenses in the LhARA Stage 1 beam-line and to investigate their capabilities in comparison to solenoid magnets, particle-tracking simulations of a realistic

3-D laser-generated proton beam were performed. The results of tracking protons through the LhARA Stage 1 beam-line with plasma lenses or solenoid magnets are compared considering two aspects. Firstly, the ability to capture protons based on their kinetic energy and initial divergence is investigated. Secondly, a comparison is presented in terms of the efficiency with which protons within the nominal energy range are selected from the wide energy spectrum and transported to the end station.

This chapter is divided into four parts, dedicated to the methods of modelling plasma lenses in a particle-tracking code (section 7.1), the simulation results for beam capture (section 7.2) and for energy selection (section 7.3), with a concluding discussion of the main results (section 7.4). The simulations presented in this chapter were performed with GPT (General Particle Tracer v2.8 [199]) where space charge was considered and with BDSIM (v1.6.0 [141]) for Monte Carlo proton tracking with no space charge.

7.1   Models of a plasma lens for particle tracking

Non-neutral plasma lenses are a novel proposal for use as proton and ion focusing devices. The current beam-tracking codes do not provide built-in routines to model the passage of the particles through the plasma lenses. Three possible solutions are presented here starting from the faster, less complex ones (replacing a lens by an equivalent solenoid or using a hard-edge uniform-plasma field map) and introducing a more realistic solution (2-D cylindrically symmetric field map generated from a fully thermalised non-uniform plasma).

7.1.1   Equivalence to solenoids

While the baseline design [270] of the Stage 1 of LhARA employs a series of five plasma lenses, an alternative beam-line design exists in which normal-conducting solenoids replace the plasma lenses. Furthermore, the particle-tracking studies required to optimise the transport of the ions from the source to the end stations require a model of the five focusing elements that allow the simulations to run relatively fast. This is the case if the focusing elements are considered to be solenoids.

The focusing elements in the LhARA beam-lines have been, so far, modelled as solenoids in order to allow for rapid tracking of the ions and, hence, efficient optimisation and studies of the beam-transport capabilities. It is argued here that the beam-transport capabilities of the LhARA beam-lines are highly comparable when using either solenoids or plasma lenses as the focusing elements. This equivalence is ensured by choosing the magnetic flux density of one solenoid and the electron density of the equivalent plasma lens such that the focusing strength is identical.

Firstly, as has been argued before (see sec. 1.3.4 and fig. 1.7), the chromatic dependence of the focusing strength on the beam energy arising from a solenoid and a plasma lens are almost the same, with identical focal lengths at the nominal energy. Thus, the energy collimation scheme that was developed for LhARA based on plasma lenses is expected to show a similar capability to remove the protons outside of the nominal energy range even when the lenses are replaced by solenoids. This is a consequence of the relatively low nominal capture
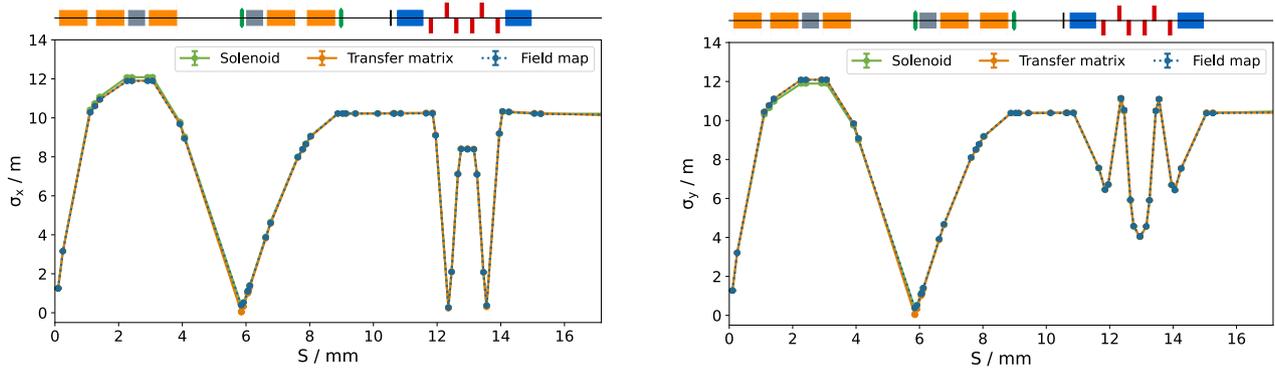
Figure 7.1: Comparison of the transverse beam size (for nominal beam energy) along the Stage 1 of LhARA for two distinct focusing elements: solenoids and plasma lenses. Each of the five plasma lenses was modelled either as a transfer matrix or with the help of an electrostatic field map. The focal length of each focusing element is identical irrespective of how the element is modelled.

energy of $15\,\mathrm{MeV}$.

Secondly, fig. 7.1 shows good agreement from particle-tracking simulations of protons transported through the Stage 1 beam-line using either solenoids or plasma lenses. The transverse beam sizes shown in fig. 7.1 were found to be similar for three different means of modelling the focusing elements. Protons with kinetic energy in the range $15\,\mathrm{MeV}\pm2\%$ were tracked with BDSIM to compare the three models. Solenoids are available in BDSIM as a predefined component. The plasma lenses were modelled either as a finite-length transfer matrix applied at the middle of the element or using electrostatic field maps. Both the transfer matrices and the field maps were generated on the assumption of an ideal plasma distribution with uniform density and hard edges. All the collimators were opened for the beam-optics study presented here.

Figure 7.1 also depicts the main actions of the beam-line on the proton beam. The first three focusing elements capture the divergent proton beam and re-focus it to a point just before the six metre mark. The last two focusing components take the beam from point to parallel and match it into the vertical arc with one-to-one optics. A highly parallel beam reaches the end-station with negligible divergence angle in both transverse planes.

A finite-length transfer matrix provides a simple model of a plasma lens, is relatively quick to generate, and allows particle-tracking codes to transport particles fast from the entry to the exit planes of the lens. However, a fundamental drawback of the transfer matrix is that the terms of the matrix are fixed, and, thus, not energy dependent. In reality, ions with different kinetic energy experience different focusing strength when propagating through the plasma lens. Since an electron plasma lens is not a typical component of particle accelerators, current particle tracking codes do not include proper routines that can change the transfer matrix of the lens based on the individual energy of each particle passing through the lens. Nonetheless, the transfer matrix represents a simple solution to model a plasma lens when the energy spread of the beam is not large.

Furthermore, the transfer matrix only allows particles to be tracked inside a uniform plasma without the option to model the focusing field inside the lens, but outside of the plasma. An alternative and more realistic model of a plasma lens is afforded based on electromagnetic field maps.

7.1.2   Field map for an ideal lens

Electrostatic or electromagnetic field maps represent a more realistic model of a plasma lens. Since the particles in the beam are transported through the lens by integrating the equation of motion, each particle experiences a focusing strength that depends on its kinetic energy. Thus, modelling plasma lenses with the help of field maps is particularly valuable when studying or optimising energy-collimation schemes or when tracking beams with large energy spread. Moreover, a field map can be generated to include the space radially outside of the plasma, but still inside the vacuum vessel of the lens. Hence, particles from the beam propagating near the edge or outside of the plasma experience a more realistic focusing force if such a field map is employed. Even though the radially focusing force outside of the plasma is non-linear and falls as $1/r$, the protons that experience this weaker focusing may still be transported downstream of the lens without hitting the inner walls of the electrodes inside the lens. As a result, a more realistic transport efficiency can be extracted from simulations using field maps in comparison to transfer matrices or solenoids for which the particles in the beam propagating outside of the linear focusing region are removed from the beam.

The simplest method to generate an idealised field map for the plasma lens is based on the assumption of a spatially uniform space-charge distribution. Thus, the plasma is considered to have a fixed nominal particle number density inside the active volume of the lens and zero particle density outside. The transition of the particle density between the volume occupied by the plasma and the vacuum is simply a step function. In other words, the plasma is modelled with hard edges both radially and along the beam axis.

A second simplifying assumption is that the overall effect of the axial component of the electric field of the lens can be neglected. The assumption is justified by the fact that ions experience a negligible net change in their longitudinal kinetic energy after passing through a long electron plasma. For example, the maximum acceleration or deceleration that ions can experience while propagating through a incompletely filled lens is given by the depth of the axial potential well generated by the electrodes of the lens. These electrodes are expected to be biased to voltages below $100\,\mathrm{kV}$. Hence, the change in the kinetic energy is expected to be smaller than $100\,\mathrm{keV}$ ($<1\%$ of the nominal energy). Given all of the above, the lens is modelled as an infinitely long column of space charge with the radial component of the electric field obtained from Poisson's equation as

$$E_r(r) = -\frac{en_e}{2\epsilon_0} \times \begin{cases} r, & r \le r_p \\ \frac{r_p^2}{r}, & r_p < r < r_w \end{cases} \tag{7.1}$$

where n_e, r_p are the nominal plasma density and radius, r_w is the inner radius of the lens electrodes, and r is the radial distance from the axis of symmetry of the lens. A 3-D Cartesian field map is then generated from n_e by specifying the electrostatic field on a regular grid of points (x, y, z) as

$$\mathbf{E}(x, y, z) = E_r \frac{x}{r} \cdot \hat{\mathbf{x}} + E_r \frac{y}{r} \cdot \hat{\mathbf{y}} + 0 \cdot \hat{\mathbf{z}}. \tag{7.2}$$

In the case of the equivalent solenoid, the field map can be generated from the nominal magnetic flux density, B_{sol}, via the conversion

$$-\frac{en_e}{2\epsilon_0} = -\frac{Ze}{4m_{\text{ion}}\gamma_0}B_{\text{sol}}^2, \tag{7.3}$$

where $Z, m_{\text{ion}}, \gamma_0$ are the atomic number, mass, and relativistic Lorentz factor of the ion species in the beam, respectively. For a more precise model of the lens, a magnetostatic field map can be superimposed on top of the electrostatic one to include the marginal focusing effect generated by the current coil that typically encapsulates the lens electrodes. For the particle-tracking studies presented in this work, field maps for ideal plasma lenses were successfully used both with BDSIM and GPT.

7.1.3 Field map for a thermalised plasma

The hard-edge, uniform-density model of a plasma lens is a practical, first-order approximation of the actual typical spatial distribution of the space charge inside the lens. The corresponding electrostatic field map can be generated quickly for a given plasma density. To produce a more realistic field map, the model must include edge effects and allow for spatial variations of the plasma density.

The approach developed in this work to obtain more realistic field maps for the plasma lenses of interest in LhARA is based on the assumption that, in a practical lens, the electron plasma reaches a state that is mostly dictated by global thermal equilibrium before the beam passes through the lens. Equally important, the plasma is considered to be confined by a uniform magnetic field and the interaction with the background gas is neglected. Under these assumptions, the spatial distribution of the confined plasma is described analytically by a pair of self-consistent Poisson-Boltzmann equations (see section 3.1.2 and eqs. (3.16) and (3.18)). Several well-established numerical algorithms are available to solve the equation

$$\nabla^2\phi = \frac{e}{\epsilon_0}n_e(r, z) = f(r, \phi), \tag{7.4}$$

and obtain a spatial distribution of space charge $en_e(r, z)$ or, equivalently, a field map of the electrostatic potential, $\phi(r, z)$, for the entire volume of the plasma lens.

Any chosen numerical procedure typically requires a significant computation time, longer than the one necessary to generate the idealised hard-edge field maps described before. In a state close to global thermal equilibrium, the charge distribution has a uniform density that falls both radially and longitudinally near the edges of the plasma over a distance of several Debye lengths. For a cold plasma, the size of the region over which the density changes from the peak value to approximately zero is, thus, much smaller than the radius or the length of the active region of the plasma lens. Therefore, the hard-edge model is still a good approximation for a lens consisting of a cold plasma in global thermal equilibrium. One advantage of the more elaborate model of the plasma is that it allows for focusing aberrations to be reproduced and studied in particle-tracking simulations. Furthermore, it represents a more realistic representation of a plasma lens if the Debye length is expected to be a considerable fraction of another dimension of interest such as the beam size or the plasma radius.

In section 3.1.2, the dependence of the azimuthal rigid-rotor angular velocity of the electrons on the radial density profile of the plasma was studied (see fig 3.4). This is a result of the electrons adapting their position in the equilibrium state to cancel out any externally imposed electrostatic field and to counter-balance the centrifugal radial force. However, in an experiment, it is the radial profile of the plasma column that is most commonly measured, and not the angular velocity ω_{re}. Similarly, when defining the requirements for a plasma lens, one important specification is the radius of the plasma. Thus, it is convenient to introduce the quantity r_p which is defined as the radius in the mid-plane, $z = 0$, at which the plasma density has fallen to half of its central value, $\hat{n}_e$. The source term of the Poisson equation (3.18) can, then, be rewritten as [271]

$$f(r, \phi) = \frac{e\hat{n}_e}{\epsilon_0} \exp\left[\frac{e}{k_B T_e}(\phi - \phi_{00}) - \alpha r^2\right],\tag{7.5}$$

with the constant α readjusted at each step of the iteration according to the formula

$$\alpha = \frac{e}{k_B T_e r_p^2}\left(\phi(r_p, 0) - \phi(0, 0)\right) + \frac{\ln 2}{r_p^2}.\tag{7.6}$$

The cylindrical geometry used for computing the plasma equilibrium has a radial extent $r_{\max}$ and a length $z_{\max}$, and is subdivided into rectangular cells with sides of length Δr and Δz. For accurate results, the size of the cell in each of the two dimensions has to be less than or equal to half the Debye length [271]. Consequently, the overall computation time for a given tolerance in the final solutions increases for a plasma with lower temperature or higher density. The plasma is assumed to be axisymmetric and the confining magnetic field uniform and parallel to the axis of the trap. Two ring-shaped end electrodes create the axial confining potential. The computation region was chosen for a plasma state with reflection symmetry about $z = 0$. Under these conditions, the corresponding boundary conditions on ϕ are

$$\frac{\partial \phi}{\partial r} = 0 \ \text{ at } \ r = 0,\tag{7.7}$$

$$\frac{\partial \phi}{\partial z} = 0 \ \text{ at } \ z = 0 \text{ and } z = z_{\max}.\tag{7.8}$$

Also, at the conducting wall $\phi = \phi(a, z)$ is zero except at the location of the end-cap electrode used to confine the plasma longitudinally. Given the reflection symmetry of the plasma column in the radial and axial mid-planes, the plasma density is computed only in a quarter of the full volume of the lens. $z_{\max}$ is taken as the transverse plane in the middle of one of the end-cap electrodes.

The iterative algorithm chosen here is successive over-relaxation (SOR) with finite-differences and Newton's method (Newton SOR). This particular iteration scheme was chosen for its simplicity and effectiveness in the present problem [271, 272]. The computation volume was divided into a regular grid and the derivatives were evaluated by the central-difference method. The solutions are found by relaxation from an arbitrary initial potential $\phi^{(0)}(r, z)$ through progressive iterations over the entire grid. During each iteration, a new electrostatic

potential is obtained at each grid point through

$$\phi^{(l+1)}(r,z) = (1-\omega)\phi^{(l)}(r,z) + \frac{\omega}{2}\left(\frac{1}{\Delta r^2} + \frac{1}{\Delta z^2} + 2\left.\frac{\partial f}{\partial \phi^{(l)}}\right|_{r,z}\right)^{-1} \times \tag{7.9}$$

$$\times \left[\frac{(1+\Delta r/2r)\phi^{(l)}(r+\Delta r,z) + (1-\Delta r/2r)\phi^{(l+1)}(r-\Delta r,z)}{\Delta r^2}\right. \tag{7.10}$$

$$+ \frac{\phi^{(l)}(r,z+\Delta z) + \phi^{(l+1)}(r,z-\Delta z)}{\Delta z^2} \tag{7.11}$$

$$\left.- f\left(r,\phi^{(l)}(r,z)\right) + \left.\frac{\partial f}{\partial \phi^{(l)}}\right|_{r,z}\phi^{(l)}(r,z)\right], \tag{7.12}$$

where $1 < \omega < 2$ is a mixing constant. A good estimate for the best value of ω is the value used to solve Poisson's equation with homogeneous Dirichlet boundary conditions on a rectangular $n_r \times n_z$ grid [273]:

$$\omega = \frac{2}{\sqrt{1-\rho^2}} \tag{7.13}$$

$$\rho = \frac{\Delta z^2 \cos(\pi/n_r) + \Delta r^2 \cos(\pi/n_z)}{\Delta z^2 + \Delta r^2}. \tag{7.14}$$

The initial potential $\phi^{(0)}(r,z)$ was typically chosen as the electrostatic potential of the trap in the absence of plasma in order to start with an initial guess that is closer to a convergent solution and, thus, speed up the calculation. The right-hand side of equation (7.12) contains a mix of iteration levels (both l and $l+1$) to make use of the updated values of ϕ as soon as they become available [271]. Furthermore, the $\partial f/\partial \phi$ term improves the algorithm for nonlinear problems by evaluating $f(r,\phi)$ at the $(l+1)$ iteration level. The potential ϕ is updated starting from the lower left-hand corner of the grid and the calculation proceeds radially first. The iterations are continued until the solution converges within the desired tolerance.

The result of the numerical procedure outlined above is a pair of two 2-D spatial distributions: the charge or particle density, and the electrostatic potential. Firstly, the charge density is useful for evaluating the linearity of the focusing force inside the lens based on the uniformity of the plasma in the region through which the ion beam passes. Secondly, the electrostatic potential can be converted into an electrostatic field map or used as it is as input for particle tracking programs.

Figure 7.2 shows an example of the electron number density obtained with the numerical procedure described above for the second plasma lens in Stage 1 of LhARA. Only one of the ends of the plasma column is shown. The plasma temperature was chosen to be relatively large ($10\,\text{eV}$) in order to speed up the computation and to enlarge the size of the region over which the plasma density rapidly drops from the nominal value in the centre of the lens. The corresponding Debye length is $0.87\,\text{mm}$ and the grid size was set to $0.2\,\text{mm}$. The electrons redistribute themselves such that the equipotential lines inside the plasma become parallel to the symmetry axis of the trap. The length of the plasma column measured near the axis approximately matches the length of the central grounded electrode of the lens. In order to achieve that, the end-cap electrode is biased at a voltage equal to approximately twice the self space-charge potential of the plasma on the axis.
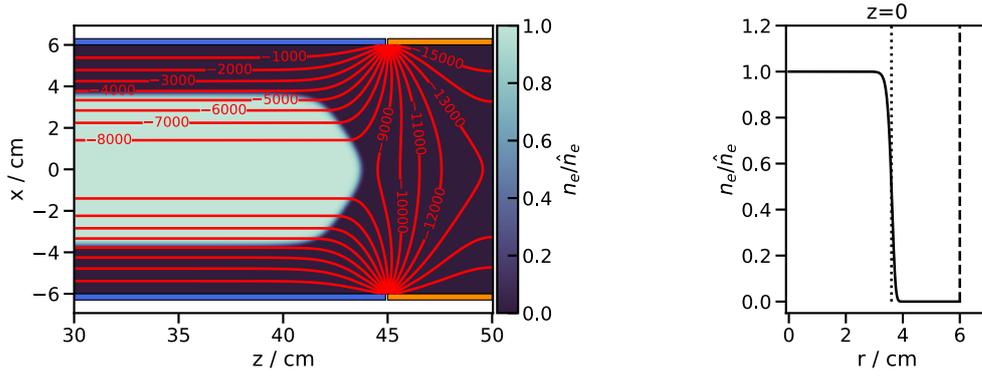
Figure 7.2: **Left:** The electron number density, n_e, near the end of a plasma column trapped in a Penning-Malmberg trap (longitudinal cross-section). The distribution is calculated numerically by solving the self-consistent pair of Poisson-Boltzmann equations (7.4) and (7.5) . The central electrode (blue) is grounded and the end-cap electrode (orange) has a bias potential of $16\,\mathrm{kV}$. The peak plasma density is $\hat{n}_e = 7.2 \times 10^{14}\,\mathrm{m}^{-3}$ and the electron temperature $T_e = 10\,\mathrm{eV}$. The red curves mark the equipotential lines and the labels indicate the corresponding potential in kV. **Right:** The radial density profile of the plasma in the mid-plane of the lens. The plasma density drops to a half of the peak value at $r_p = 3.6\,\mathrm{cm}$ (dotted line) and is zero at the conducting wall ($r_w = 6\,\mathrm{cm}$, dashed line).

In the particular example shown in fig. 7.2, the plasma column is by approximately 5% shorter at $r = r_p$ compared to $r = 0$. Thus, the protons that propagate at large excursion from the beam axis experience a focusing strength as low as 95% of the strength experienced by protons that travel close to the axis. The shape of the plasma at the two ends is dictated mainly by the bias voltage applied to the end-cap electrodes in relation to the self space-charge potential of the plasma on the axis.

The numerical procedure outlined here was employed to generate a separate field map for the first three plasma lenses in Stage 1 of LhARA. The central density for each lens is specified in table 7.2 and a plasma temperature of $1\,\mathrm{eV}$ was used throughout all calculations. The bias voltage of the end-cap electrodes was adjusted for each lens to ensure that the length of the plasma on axis was approximately equal to the nominal optical length of $0.857\,\mathrm{m}$. The three more realistic field maps of the electrostatic potential $\phi(r, z)$ were successfully used in GPT for tracking protons through the first $6\,\mathrm{m}$ of the LhARA Stage 1 beam-line with space charge.

While resulting in more realistic field maps for the plasma lenses, the numerical procedure of calculating the distribution of the plasma in global thermal equilibrium requires significantly more computation time in comparison to the generation of hard-edge field maps from equation (7.1). A proton sample from a laser-generated beam was tracked with GPT through the first three lenses in LhARA with full space-charge effects to compare the impact of the shape of the plasma edge on the transport of the beam. Figure 7.3 shows the beam size tracked through the field maps of the first three plasma lenses for the hard-edge analytical model and the global thermal equilibrium numerical model of a lens. Accordingly, there is a marginal difference that results from the two models, so that the non-linear focusing generated by the shape of the plasma column at the outer edge can be neglected during beam-transport simulations.
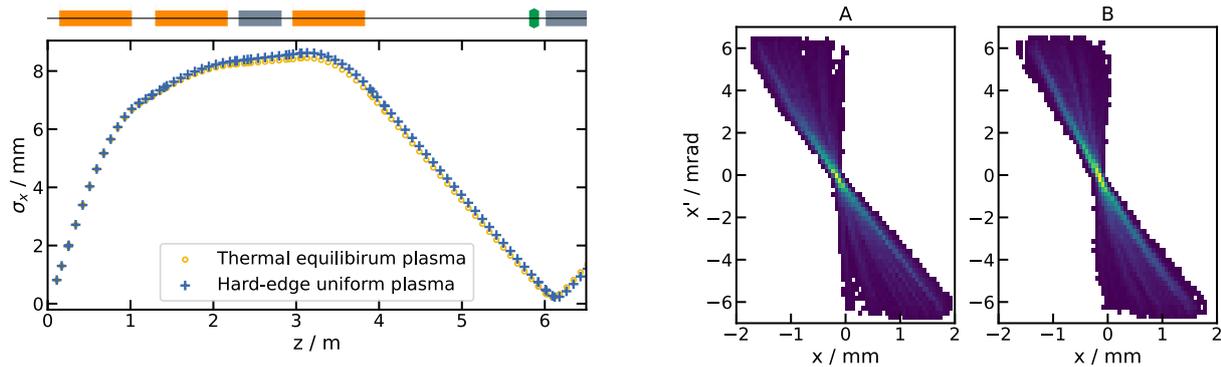
Figure 7.3: **Left:** Beam size along the front-end of the LhARA Stage 1 beam-line for protons tracked with GPT. The plasma lenses were modelled as field maps obtained under the assumption of a hard-edge uniform plasma (see eq. 7.1) or of a plasma in global thermal equilibrium (see section 7.1.3). **Right:** Phase space distribution close to the focal point downstream of the third lens for the two different models of the plasma lenses.

7.2 Beam capture

The beam optics in the Stage 1 of LhARA and the capability of the beam-line to deliver a wide range of beam sizes to the low-energy *in-vitro* end station have been validated prior to this work under a couple of scenarios. The studies were carried out with ideal Gaussian proton and carbon beams at the source [9] and with a 3-D proton distribution sampled from a 2-D PIC simulation of the laser interacting with a thin-foil target [247]. As part of the work towards the conceptual design report for LhARA, 3-D PIC simulations were carried out to generate a realistic proton spectrum that has the main features expected to be afforded by the source planned for LhARA. Thus, the aim of the work presented in this section is to corroborate the previously simulated beam-transport capabilities of Stage 1 of LhARA with a realistic 3-D proton distribution at the source. While a more comprehensive description of the 3-D PIC simulation that generated the proton distribution can be found elsewhere [274, 275], a summary of the relevant parameters is provided below.

7.2.1 Proton distribution at the source

The technology required for the laser-driven ion source for LhARA has started to be tested and developed as part of the preliminary R&D work for the facility [275]. Preliminary simulations of the interaction between a laser and a thin solid target have been performed and analysed to predict an optimised proton source for a 100+ TW laser system. The output of the preliminary simulations (courtesy of E. Boella, Lancaster University) were employed in this work for modelling the LhARA Stage 1 beam-line with more realistic beam parameters. In parallel to the simulations, an experimental campaign has started at the SCAPA laser facility [276] with the aims to produce initial parameter scans of the proton and ion beam properties, develop high-repetition rate targetry, and quantify debris production [275]. As the experiments at SCAPA will provide data sets suitable for benchmarking the simulation work, the preliminary simulations focused on modelling one of the SCAPA lasers which has similar pulse specifications to the proposed LhARA ion source.
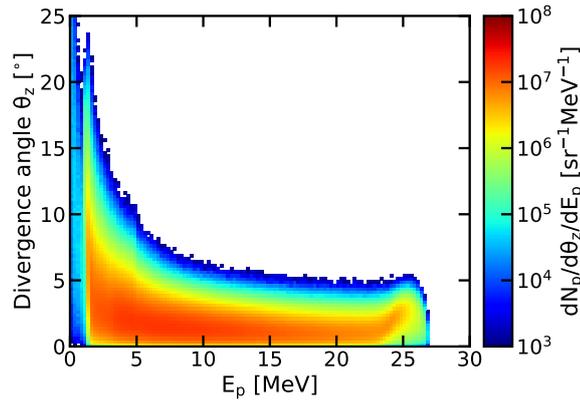
Figure 7.4: Divergence angle, θ_z, of the protons as a function of their kinetic energy, E_p, at the laser-driven source. Distribution obtained from a PIC simulation of the interaction between a laser and a solid tape target [274, 275].

The 3-D particle-in-cell simulations explored the interaction of a Ti:sapphire laser (central laser wavelength $\lambda_0 = 800$ nm) with aluminium foils. The main parameters of the particular simulation of which results are referred to in this work are given in table 7.1. The target was modelled as a pre-formed plasma composed of Al^{3+} ions and electrons with density $70\,n_c \sim 10^{23}\,\mathrm{cm^{-3}}$ (here n_c is the critical density corresponding to the laser frequency). A thin layer of H^+ ions was inserted on the back of the target to mimic contaminants naturally present on the back surface of solid targets. The presence of the pre-plasma is well-known to enable more efficient electron heating mechanisms in the target which translate into higher proton energies. For LhARA, it is desirable to achieve an energy cut-off well above $15\,\mathrm{MeV}$ to increase the amount of charge produced within the nominal energy band and to use a part of the spectrum that is less affected by experimental fluctuations.

Table 7.1: Parameters used in the 3-D PIC simulation of a SCAPA-like laser [276] interacting with a thin-foil solid target [274].

Parameter	Value	Unit
Laser intensity	9×10^{20}	$\mathrm{W/cm^2}$
Laser wavelength	800	nm
Laser pulse length	25	fs
Normalised laser amplitude	20.52	
Laser spot size	1.5	µm
Target thickness	6	µm

The proton distribution was extracted from the PIC simulation at the last available time step, when the proton cut-off energy does not change anymore. A proton spectrum extending beyond $20\,\mathrm{MeV}$ in kinetic energy was achieved in the simulation with a target thickness of $6\,\mathrm{µm}$ and laser intensity $\lesssim 10^{21}\,\mathrm{W/cm^2}$. The energy-angle spectrum of the protons that are accelerated from the source is shown in fig. 7.4. The typical feature of protons generated by the TNSA mechanism are observable in the spectrum: sharp cut-off in kinetic energy
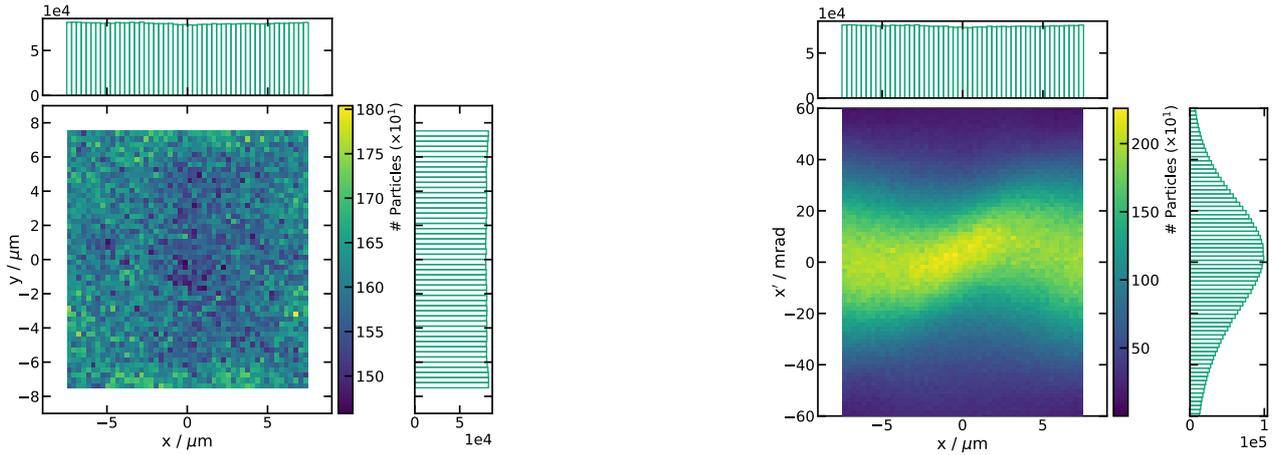
Figure 7.5: Distributions of the protons at the laser-driven source in the x–y (left) and x–x' (right) projections of the phase space. The limits of the space occupied by the protons in the x–y plane are dictated by the size of the simulation box used to obtain the distributions via the PIC method.

at about $25\,\text{MeV}$ and more energetic protons are produced at lower solid angles. In this case, the divergence angle is calculated as $\theta_z = \tan^{-1}(p_r/p_z)$. The standard deviation of the divergence angle around $15\,\text{MeV}$ is approximately $1°$. The total charge of the protons over the full energy spectrum is $2.28\,\text{nC}$, while the nominal energy band of $15\,\text{MeV}\pm2\%$ contains $39\,\text{pC}$.

The size of the area of the target from which protons are accelerated is limited by the overall size of the grid defined in the PIC simulation. Since a larger grid leads to a significantly higher computation cost, the transverse size of the proton beam at the source is $15\,\mu\text{m}\times15\,\mu\text{m}$. Figure 7.5 shows the transverse phase space of the protons at the source as generated by the PIC simulation of the laser interacting with the target. A small deviation from zero is observed for the average divergence angle in the transverse direction. The RMS emittance at the source is $6.6 \times 10^{-8}\,\text{m}\,\text{rad}$.

7.2.2 Capture and transport efficiency

The first two plasma lenses or solenoids downstream of the laser-driven source were designed in LhARA to capture as many protons as achievable from the particle flux generated at the target. The two focusing components bring the beam from point to parallel. A third lens or solenoid re-focuses the beam at the location of the energy-selection collimator. It is this collimator and a vacuum nozzle attached at the interface between the vessel of the target and that of the first lens that remove a significant number of particles from the beam generated at the source. Most of the ions that are removed from the beam have the kinetic energy outside of the nominal range, but some ions with the desired energy are also lost in the energy-selection procedure.

While a preliminary configuration of the vacuum nozzle and the energy-selection collimator was produced for the baseline design of LhARA [270], some modifications are expected to appear as a result of the practical considerations and the R&D work for building the LhARA facility. For example, the geometry of the vacuum

nozzle may suffer modifications during the design of the pumping systems once more data becomes available on the vacuum conditions achievable both around the laser-driven target and in a practical plasma lens. Furthermore, an RF cavity is planned to provide manipulation of the longitudinal phase space of the ions such that a larger fraction of the beam from the source passes through the energy-selection collimator.

Given the currently ongoing development of the design of the LhARA beam-lines, a direct comparison is presented here between the capture and transport efficiency of the equivalent solenoids versus the first three plasma lenses in the absence of any beam collimation. The protons were tracked up to the position of the energy-selection collimator to identify the fraction of the initial particle distribution from the source that is effectively captured and transported to this location. Thus, the results constitute an evaluation of the upper limit of the total number of protons that can be captured and transported from the source with the approach chosen for LhARA with a proton distribution that is similar to the ones generated by the TNSA mechanism. Alternatively, the following investigation shows the benefit of using plasma lenses over solenoid magnets for capturing protons from a laser-driven source in applications where the wide energy spread is not an issue and the highest achievable beam current is desirable.

The protons were tracked with GPT from the source to the energy collimator at approximately 6 m downstream of the source. For the first 5 cm, the space charge of the proton beam was considered to be fully compensated by the co-propagating electrons that are simultaneously accelerated from within the target. Starting from the plane at 5 cm and going downstream to the energy-selection collimator, a calculation of the full space charge of the beam was included in GPT. The solenoids were modelled with the `bzsolenoid` element available in GPT. With this element, the field of the magnet is generated by analytical expressions based on modelling the solenoid as a circular current sheet [199].

The plasma lenses were modelled as 2-D cylindrically symmetric field maps. GPT allows the electrostatic potential map to be provided as an input. Subsequently, the components of the electrostatic field are calculated using bilinear interpolation from the potential. It is known that discontinuities in the electric field appear as a result of the linear interpolation of the potential. However, using a map of the potential for the lens is favourable here since the numerical procedure for calculating the distribution of the plasma in a lens directly produces the spatial map of the potential as well. The field maps were calculated with the numerical procedure described in section 7.1 for a plasma in global thermal equilibrium. The nominal parameters of the focusing components are specified in table 7.2. Moreover, the aperture of the solenoids (3.6 cm radius) was set to match the radius of the plasma in the lenses. For a more realistic model of the lens, the inner conducting wall of the lens was placed at an appropriate distance away from the surface of the plasma, at a radius of 6 cm. As a result, the lenses have a larger geometrical acceptance compared to the solenoids, but the focusing force outside of the plasma is non-linear and decreases from the surface of the plasma to the inner wall of the lens.

Figure 7.6 shows the capture and transport efficiency through the first three focusing components of LhARA for both plasma-lens and solenoid focusing. The percentage of the protons that survive from the source to the energy-selection collimator is shown in bins of kinetic energy and divergence angle ($\theta_z = \arctan(p_r/p_z)$ where
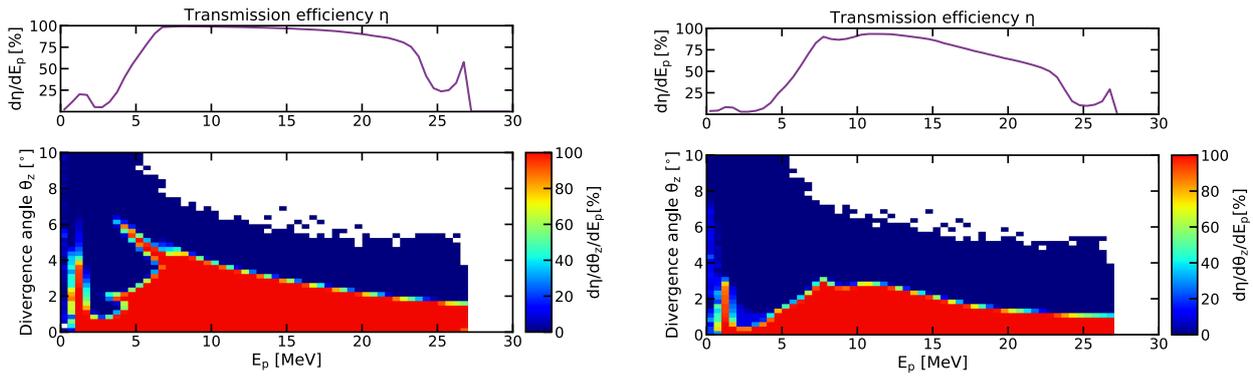
Figure 7.6: Percentage of the protons transported from the laser-driven source to the entry plane of the energy collimator in the Stage 1 of LhARA with the use of plasma lenses (left) and solenoids (right). The colour density indicates the transport efficiency. The curves on the top show the transmission coefficient integrated over the full solid angle as a function of the kinetic energy of the protons, E_p.

p_i is a component of the proton's momentum). The transmission efficiency is also integrated over the entire solid angle in fig. 7.6 to show the percentage of the protons from the source effectively transported as a function of their kinetic energy. In both the case of plasma lenses and of the solenoids, most of the protons with kinetic energy below about $5\,\mathrm{MeV}$ are lost due to their faster transverse expansion caused by the space charge. A pseudo-flat plateau in the transmission is observed between 10 and $15\,\mathrm{MeV}$. For energies higher than $15\,\mathrm{MeV}$, the transmission starts to decrease faster in the solenoid-focusing case.

Regarding the angle measured at the source of the protons that are effectively transported, there is a sharp upper limit for each energy band. The protons with kinetic energy in the range 10–$15\,\mathrm{MeV}$ that are captured and transported downstream emerge from the laser-driven target at angles below approximately $4°$ if plasma lenses are used. This limit decreases to approximately $3°$ in the case of using solenoid magnets. Hence, the percentage of the protons transmitted around $15\,\mathrm{MeV}$ is close to 100% in the case of plasma focusing and around 85% for magnetic focusing. The difference arises from the protons that emerge from the target at larger solid angles. Ideally, the angular acceptance and the entry diameter of the vacuum nozzle placed before the first plasma lens should be matched with the maximum angle with which the protons that can be captured around the reference energy emerge from the source.

For an application where a larger total beam current or charge per bunch is more preferable than a small energy spread, the transmission coefficient can be integrated over the entire energy spectrum. Equivalently, figure 7.7 shows the cumulative beam loss as a fraction of the total number of protons generated at the source. For the full beam with a wide energy spectrum, the plasma lenses have a higher capture efficiency as they lead to lower beam losses by about 15% of the initial number of protons at the exit of the third focusing component. For a much lower energy spread centred around $15\,\mathrm{MeV}$, the difference in the beam loss becomes less significant with approximately 10% of the protons transmitted by the plasma lenses in addition to the particles transmitted
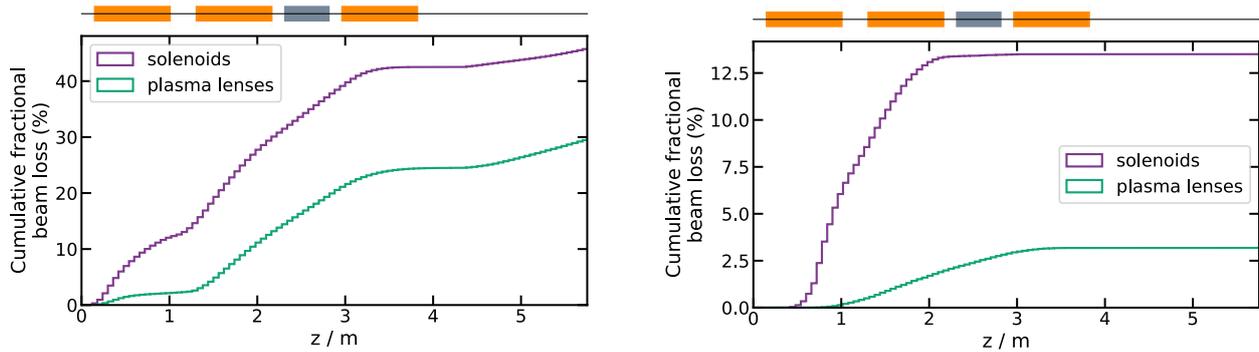
Figure 7.7: Cumulative percentage of protons that are lost from the source to the energy collimator with the use of solenoids compared to the use of plasma lenses. The beam loss is shown for the full energy spectrum of the particles from the source (left) and for the nominal energy interval of $15\,\mathrm{MeV}\pm2\%$ (right).
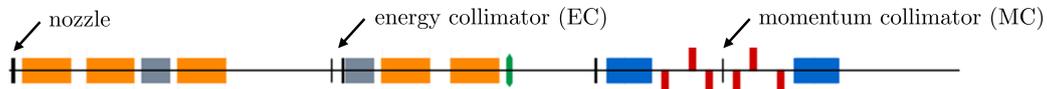
Figure 7.8: The location of the main beam collimators in the Stage 1 of LhARA. The elements of the beam-line are schematically shown as follow: the orange boxes represent solenoids, grey are RF cavities, green are octupoles, black are collimators, blue are dipoles and red are quadrupoles [9].

from the source by the solenoid magnets. In other words, 10% less protons with nominal energy from the particle flux produced at the source is transmitted by the front-end section of the LhARA Stage 1 beam-line when replacing the plasma lenses by solenoids.

7.3 Beam collimation and energy selection

Optimal beam collimation is an important requirement for the Stage 1 beam-line of LhARA. The goal is to remove from the particle distribution generated at the source the protons and ions that have large excursions from the centre of the beam or that lie outside of the nominal energy interval. Simultaneously, the procedure for beam collimation and energy selection must allow a high percentage of the particles within the nominal energy range to survive until the end station. Consequently, an energy selection scheme was developed for Stage 1 of LhARA based on a series of collimators or slits as described in the baseline design [270]. The position of the main collimators is indicated in fig. 7.8 and their role is described in the following.

The first stage of beam collimation takes place at the interface between the vacuum vessel of the laser-driven target and the vacuum vessel of the first plasma lens. The background-gas pressure desirable inside the plasma lens (10^{-9}–$10^{-8}\,\mathrm{mbar}$) is at least two or three orders of magnitude smaller than the poor vacuum conditions achievable around the target (10^{-6}–$10^{-5}\,\mathrm{mbar}$) due to ablation of the tape. As a result, a re-entrant nozzle was added to the interface to provide a differential pressure of about three orders of magnitude. The vacuum nozzle, diagrammatically shown in fig. 7.9, can be pictured as a conical frustum with an entrance radius of $2\,\mathrm{mm}$, exit
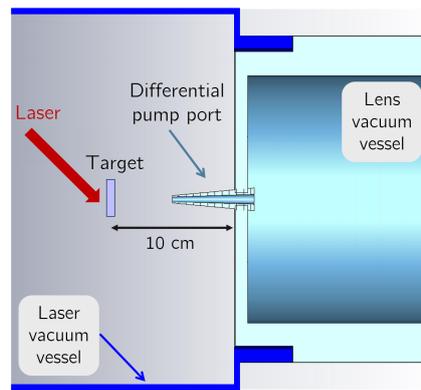
Figure 7.9: Schematic diagram of the interface between the vacuum vessel of the target and the vacuum vessel of the first plasma lens. A differential pumping port was designed [270] to provide the difference in the required background-gas pressure between the two systems.
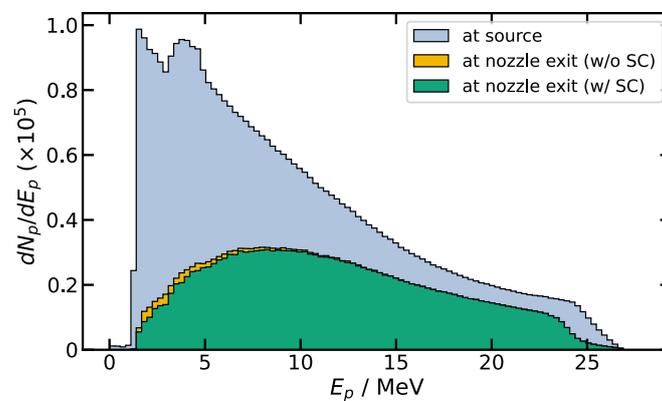
Figure 7.10: Distribution of the kinetic energy, E_p, of the protons at the source and at the exit plane of the vacuum nozzle (see fig 7.9) as simulated with GPT. The protons were tracked both with and without space-charge effects within the 5 cm-long nozzle.

radius of $2.87\,\text{mm}$, and a total length of $5\,\text{cm}$. The small entrance and exit apertures, as well as the tapering angle of the nozzle impose a geometrical acceptance on the particles that pass from the source into the first plasma lens.

The beam collimation generated by the vacuum nozzle is illustrated in fig. 7.10 for protons. The main impact of the nozzle is the absorption of the particles from the source with kinetic energy below approximately $10\,\text{MeV}$. Some of the protons within the nominal energy range around $15\,\text{MeV}$ are also lost due to their divergence being larger than the tapering angle of the nozzle. The particle distributions shown in fig. 7.10 were obtained by tracking protons from the source to the exit plane of the nozzle in GPT. Five separate beams of four million protons were sampled from the initial large proton distribution at the source generated with a PIC simulation and shown in fig. 7.4. Each beam was tracked separately and the resulting energy distributions were averaged. For the first $5\,\text{cm}$, the beam was tracked with no space charge assuming that electrons produced at the rear surface of the target co-propagate with the protons and fully compensate the space charge effects. From 5 to $10\,\text{cm}$, two cases were investigated with and without space charge. The self repulsion of the protons in the presence of space charge leads to marginally more protons being lost at the low energy end of the spectrum. While space charge
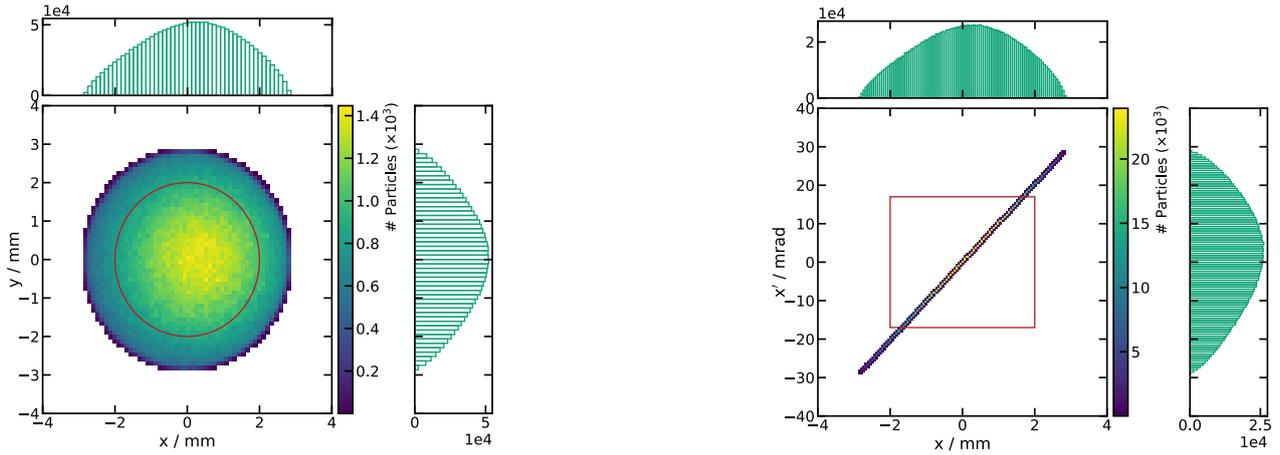
Figure 7.11: Distributions of the protons at the entry plane of the vacuum nozzle in the x–y (left) and x–x' (right) projections of the phase space. The red circle and rectangle mark the geometrical acceptance of the nozzle at the entry face.

may lead to redistribution of the protons in phase space, the overall number of protons that pass through the nozzle is mainly dictated by the acceptance of the nozzle with reduced impact of the space charge. Figure 7.11 shows the phase space of the beam at the entry plane of the nozzle compared to the geometrical acceptance.

The sampled proton beams that were tracked with GPT with space charge up to $10 \, \text{cm}$ were subsequently tracked with BDSIM through the rest of the Stage 1 beam-line to investigate the efficacy of the energy-selection scheme. No space charge was included after the first $10 \, \text{cm}$ since the large number of protons being tracked leads to large CPU time if space charge effects and field maps are simultaneously used in a simulation. Figure 7.12 shows the evolution of the transverse beam size from the exit plane of the nozzle to the *in-vitro* end station. The gradual reduction of the energy range of the protons within the transmitted beam is shown at several key locations along the beam-line in fig. 7.13.

The protons that emerge from the nozzle still have a wide energy spectrum, the full beam is relatively large when it reaches the energy collimator (EC) positioned around the $6 \, \text{m}$ mark. The large beam size is caused by the lower energy particles being over-focused and the higher energy protons not focused strongly enough by the first three solenoids. Only a small fraction of the low-energy protons are lost in the first three solenoids. The location of the energy collimator is chosen to correspond to the position of the focal point for the protons with nominal energy. As a result, the protons within the nominal energy interval pass through a small opening in the collimator. The aperture can be optimised based on the phase space, shown in fig 7.14, of the protons within the nominal energy interval at the location of the collimator. The radius of the EC collimator (circular aperture) in the baseline design is $5 \, \text{mm}$. It must be noted that the energy collimator absorbs most of the protons outside of the nominal energy range and generates an energy distribution that is peaked at $15 \, \text{MeV}$ and has wide symmetric tails around the peak. Hence, an additional collimator is required to remove the tails of the energy distribution.

The additional collimator (MC) is positioned in the middle of the vertical arc at a location with large
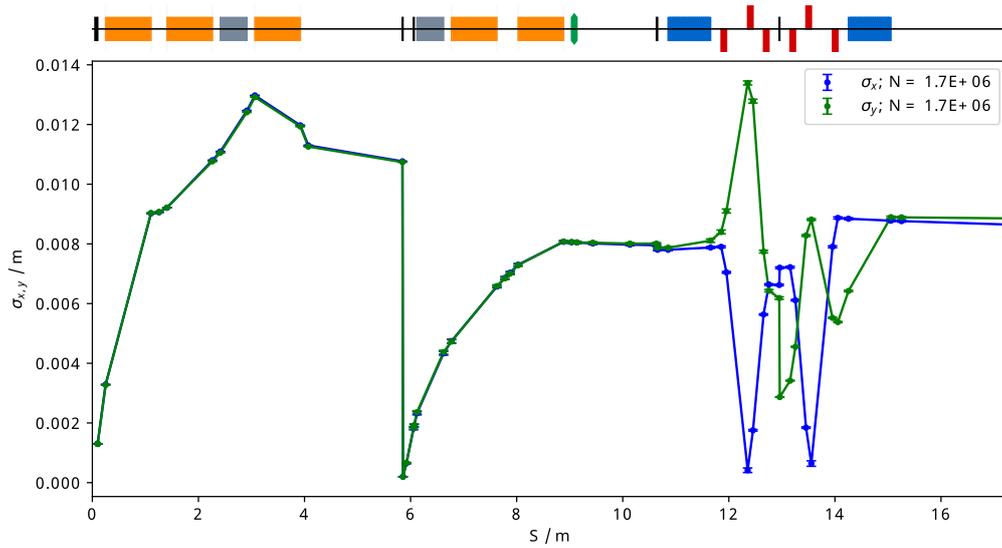
Figure 7.12: The evolution of the transverse beam size from the source to the *in-vitro* end station of LhARA obtained with a proton distribution with a large energy spread typical to laser-driven sources. In the diagram above the plot, the orange boxes represent solenoids, grey are RF cavities, green are octupoles, black are collimators, blue are dipoles and red are quadrupoles [9].
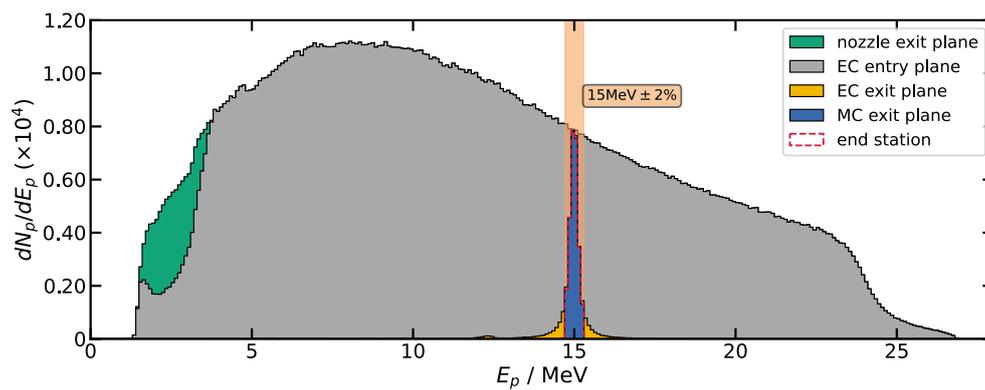
Figure 7.13: Distribution of the kinetic energy, E_p, of the protons at several key locations along the LhARA Stage 1 beam-line. The energy collimator (EC) and momentum collimator (MC) remove the protons outside of the desired energy range of $15\,\mathrm{MeV}\pm2\%$. In the simulation shown here, the focusing elements of the beam-line are solenoids.
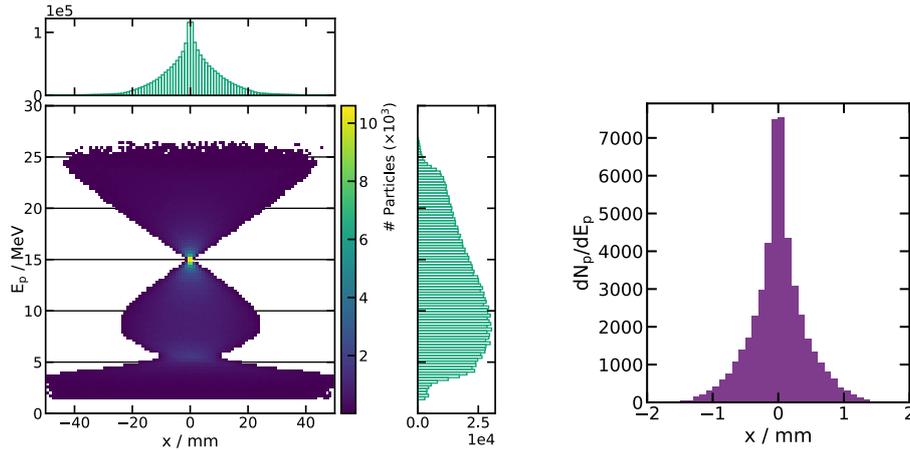
Figure 7.14: Distribution of the protons along one of the transverse directions, x, at the entry plane of the energy collimator situated after the third plasma lens in the Stage 1 of LhARA. In the figure on the left, the protons are separated by their kinetic energy, E_p, while in the figure on the right, the distribution corresponds to the protons inside the nominal energy interval of $15\,\mathrm{MeV}\pm2\%$.

dispersion in one of the transverse directions. Hence, the protons with a large momentum deviation from the nominal value exhibit a large excursion from the central beam axis. A slit placed at this location removes protons at large radii and, thus, results in 'momentum cleaning' of the beam. Figure 7.13 demonstrates that the MC slit removes the tail of the energy distribution resulting from the energy collimator. Moreover, more than 98% of the protons within the nominal energy range that arrive at the MC are transmitted downstream.

Finally, the protons that pass all the collimators are transported to the end station. The distribution of the beam at the *in-vitro* end station is shown in fig. 7.15. In the x–y plane, the observed beam asymmetry is caused by a non-zero average beam divergence at the source. Since the octupole magnets were turned off for the studies presented here, the profile of the deposited dose is circular and peaked at the centre. In the x–x' plane, the beam is observed to be highly parallel at the entrance of the end station.

The results presented in this subsection were obtained by tracking protons with a wide energy spectrum through the Stage 1 beam-line with five solenoids. As the focusing strength of a solenoid changes with the beam energy (non-relativistic case) in a comparable manner to the focusing strength of a plasma lens, a further validation is included here. The five sampled proton beams were separately tracked through the Stage 1 beam-line with field maps of plasma lenses replacing the five solenoids. Table 7.2 specifies the nominal parameters for the five focusing components in both cases. Each field map was generated starting from a uniform density plasma with hard edges and an electron density that results in an identical focusing strength as the equivalent solenoid. Since it takes more CPU time to track a large number of protons through a field map compared to the case of a solenoid, only a part of the energy spectrum at the exit of the nozzle (12–18 MeV) was considered here.

Figure 7.16 provides a comparison between the capability of the beam-line to filter the protons based on
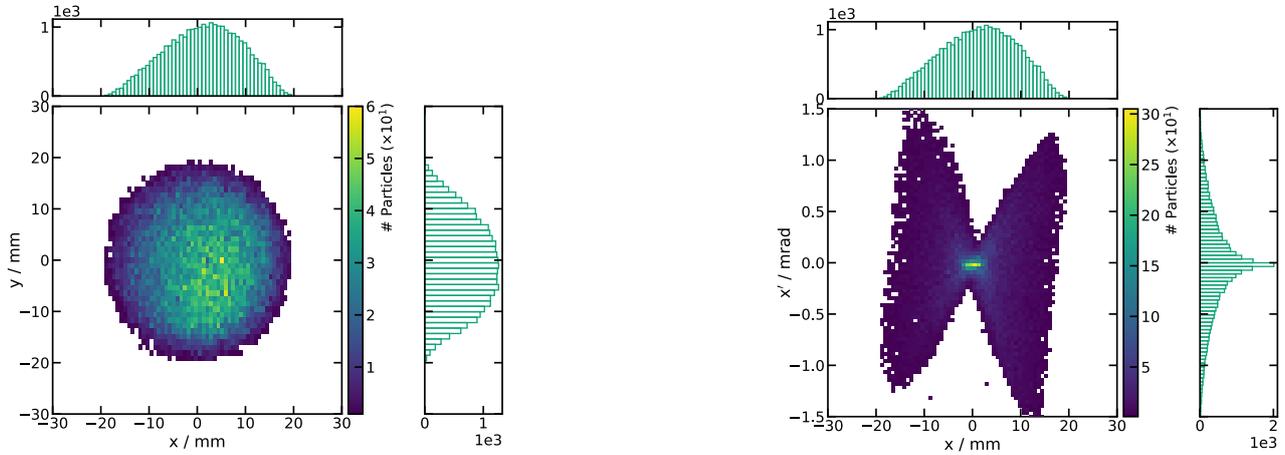
Figure 7.15: Distributions of the protons at the end station of Stage 1 of LhARA in the x–y (left) and x–x' (right) projections of the phase space.

Table 7.2: Nominal electron plasma density of the five Gabor lenses in the baseline design of Stage 1 of LhARA [270]. The table also specifies the flux density required from a solenoid with an equivalent focusing strength and identical effective length of the hard-edge focusing area (0.857 m-long).

Element	Plasma density [m^{-3}]	Magnetic flux density of equivalent solenoid [T]
Gabor lens 1	5.4×10^{15}	1.4387
Gabor lens 2	7.2×10^{14}	0.5271
Gabor lens 3	1.7×10^{15}	0.8139
Gabor lens 4	1.4×10^{15}	0.7284
Gabor lens 5	1.0×10^{15}	0.6338
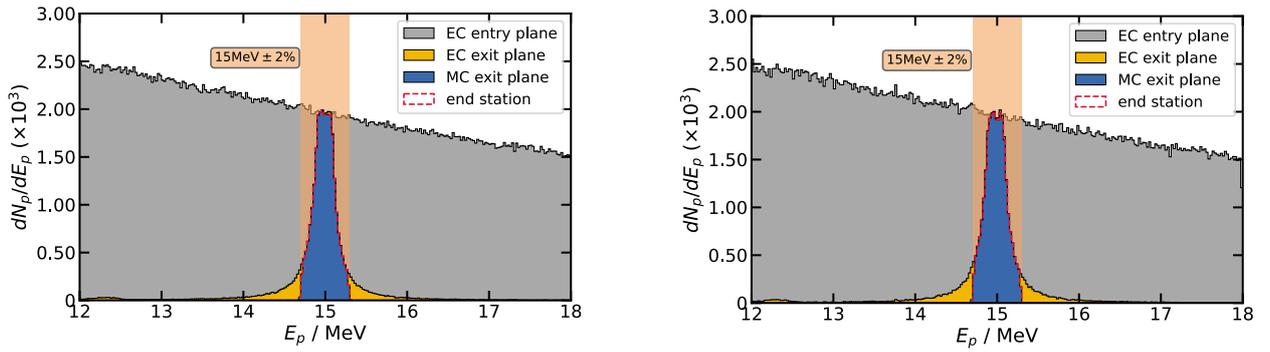
Figure 7.16: Comparison between the distributions of the kinetic energy, E_p, of protons at several key locations along the LhARA Stage 1 beam-line. The energy collimator (EC) and momentum collimator (MC) remove the protons outside of the desired energy range of $15\,\mathrm{MeV}\pm2\%$. In the simulation shown here, the focusing components of the beam-line are solenoid magnets (left) and plasma lenses (right).

their energy when solenoids or plasma lenses are used as the main focusing components. The energy profiles of the protons that reach the end station are highly comparable. Furthermore, 56% and 57% of the protons within the energy band of $15\,\mathrm{MeV}\pm2\%$ survive from the exit plane of the nozzle to the end station with the configuration of the collimators chosen for the LhARA Stage 1 baseline design [270] with solenoids and plasma lenses, respectively. Approximately 0.2% of the protons that arrive at the low-energy *in-vitro* end station are outside of the nominal energy range.

The total number of protons reaching the end station in the two cases is approximately 42% of the protons generated at the source within the nominal energy range. Overall, the superior capture efficiency of the plasma lenses demonstrated in fig. 7.6 is suppressed by the limited angular acceptance of the nozzle. In other words, the protons that emerge from the source at larger angles with respect to the beam axis and can be captured and transported by the first three plasma lenses, but not by the equivalent solenoids, are absorbed by the nozzle upstream of the capture section. As a result, similar fractions of protons within the nominal energy spread reach the end station from the source with the use of solenoids or plasma lenses. The findings presented in this chapter motivate an update of the design of the vacuum nozzle that could lead to a higher fraction of the proton flux produced at the source to be transmitted to the *in-vitro* end station.

7.4   Discussion

The comparison of the solenoid magnets to the plasma lenses for LhARA in particle-tracking simulations was motivated by the potential gain in the number of protons within the nominal energy range captured from the laser-driven source and transported to the *in-vitro* end-station with the use of lenses. Another aspect was the capability of the beam-line to select protons based on their kinetic energy. The selection efficiency was expected to be similar in the case of using plasma lenses or solenoids based on the chromatic dependence of their focusing strength on the energy of single protons in the beam.

Evaluating the performance of the plasma lenses in the LhARA Stage 1 beam-line required a method to model the transport of protons through a non-neutral electron plasma using a pair of existing particle-tracking codes. It was shown that **good agreement in the beam optics is achieved when the plasma lenses are replaced with solenoids with an identical focusing strength** as the lenses. The main advantage of using solenoids to model the plasma lenses stems from the built-in elements in both BDSIM and GPT that allow tracking particles though the field of a solenoid magnet. Replacing the lenses by solenoids in the simulations represents an effective solution during optimisation of the beam-lines when the strength of the focusing components need to be varied relatively fast. Tracking a proton beam with a small energy spread of a few percent led to results that agree well in the case of the lenses being modelled by solenoids, transfer matrices, and hard-edge uniform-plasma field maps. In addition, **minor differences were obtained in the proton beam size when using field-maps from the hard-edge model of the plasma compared to using a more realistic space charge distribution with rounded edges**. These edge effects are more significant when the diameter of the beam is comparable to the diameter of the plasma, a condition met within the third lens of the LhARA Stage 1 beam-line.

Simulations of the proton capture and transport efficiency through the first three focusing components of LhARA showed that **plasma lenses can deliver more protons to the rest of the beam-line by up to approximately 15% of the number of particles produced at the source compared to the number of protons transported by solenoids**. It must be noted that the benefit of using plasma lenses over solenoid magnets is more significant in the absence of beam collimation by the vacuum nozzle or an energy-selection slit. Consequently, the design of the vacuum nozzle for LhARA needs to be reconsidered by taking into account the transport of protons with a realistic distribution at the source characteristic to the TNSA mechanism.

In other words, the lenses have a higher capacity to capture and transport the protons that emerge from the target at larger angles and at the low and high ends of the full energy spectrum. Such a capability is desirable in an application where a high beam current is more important than a small energy spread. Alternatively, within the energy range of $15\,\mathrm{MeV}\pm2\%$, the first three plasma lenses were found to capture and transport more protons by approximately 10% of the total number of particles produced at the source compared to the number of protons transmitted by the equivalent solenoids. The first three plasma lenses were found to capture and re-focus the protons from the target that emerge inside a cone with the half-angle of approximately $4°$ and $2°$ for the kinetic energy of protons centred around $10\,\mathrm{MeV}$ and $20\,\mathrm{MeV}$, respectively.

Apart from the capability of the plasma lenses and the solenoid magnets to capture and transport the protons from the source, a set of simulations was carried out to evaluate the effectiveness of the energy collimation scheme. The start-to-end simulations tracked a realistic laser-driven proton distribution with a wide energy spread. Firstly, it was found that the low energy protons below $5\,\mathrm{MeV}$ are removed by the nozzle at the interface between the vacuum vessel of the source and that of the first plasma lens. Some protons within the nominal energy range of $15\,\mathrm{MeV}\pm2\%$ are also absorbed by the nozzle. Secondly, the energy collimator situated downstream of the third plasma lens removes most of the protons outside of the desired energy interval. The resulting energy distribution is peaked at $15\,\mathrm{MeV}$ and exhibits long symmetric tails outside of the nominal energy range.

The tails are effectively removed by the momentum-cleaning collimator situated in the middle of the vertical arc. Approximately 40% of the protons within the nominal energy range that pass the nozzle are lost in one of the other collimators before reaching the *in-vitro* end station with the specifications of the collimators in the LhARA baseline design [270].

Finally, **replacing the plasma lenses with solenoid magnets with identical focusing strength has little impact on the effectiveness of the energy selection scheme for a beam composed of only protons**. The simulations showed that the energy distribution of the protons that reach the end station are highly similar in the two cases with almost identical total number of protons successfully transported by the beam-line from start to end. As a future step in the development of the design of LhARA, it is still needed for simulations to demonstrate the capability of the beam-line to select protons or a single ion species from the multitude of particles produced at the source. The investigation of the capability of the beam-line to select a single ion type requires realistic particle distributions generated by the TNSA mechanism for ion species heavier than the proton.

For LhARA, the non-neutral electron **plasma lenses remain the baseline solution for the capture and transport of ions** from the laser-driven source. Nonetheless, replacing the lenses with solenoids was shown to provide equivalent capabilities of the Stage 1 beam-line to transmit part of the proton flux from the source to the low-energy *in-vitro* end station. Irrespective of the technology used for focusing the ions in LhARA, the **design of the vacuum nozzle at the interface between the source and the capture section of the beam-line needs to be revisited**. On the one hand, the design of the nozzle needs to be driven by the vacuum conditions achievable inside the vessel of the target and the vessel of the first focusing component, the capabilities and arrangement of the vacuum pumps, and the final design of the first focusing element in the beam-line. On the other hand, the size and divergence of protons or ions at the entrance of the nozzle should also feed into the design of the nozzle. The goal is to ensure that the aperture limitations of the nozzle do not reduce the particle flux transmitted from the source below the number of protons or ions that can be captured and transmitted by the front-end section of the beam-line in the absence of the nozzle.

In conclusion, the non-neutral electron plasma lenses and solenoid magnets represent almost identical solutions for focusing and transporting a laser-driven proton beam in LhARA Stage 1 when considering their capabilities for capturing the particles from the source and reducing the initial wide energy spectrum of the beam. According to the studies presented here, there is, nonetheless, a potential gain in the total number of protons that can be captured from the target by using plasma lenses if the interface between the source and the first lens will suffer further updates of its design. While there are other strategies to increase the number of protons per bunch delivered by the LhARA beam-line to the *in-vitro* end station, such as the use of a laser with a power greater than $100\,\mathrm{TW}$, the reconsideration of the design of the vacuum nozzle is the most cost-effective solution given that the conceptual design of LhARA is currently under development.

The final conclusions drawn from the work performed during my PhD as well as some perspectives for the future development of plasma lenses for LhARA and beyond are presented in the following chapter.

Chapter 8

Conclusions and perspectives

Despite major research efforts in the last decades, the reliable operation of a high-density non-neutral plasma lens for focusing charged particle beams has not yet been proven. Most of the attempts to create a practical lens resulted in devices that were able to produce weakly confined quasi-steady-state plasmas. The associated operating conditions of the lens most frequently resulted in a plasma regime unfavourable for beam focusing. Future advances are likely to be found in the exploitation of new insights into the limitations of the approaches explored so far and in the use of fundamentally different strategies for producing and controlling the plasma within the lens. Further research efforts are motivated by the prominent advantage that the plasma lenses present over conventional solenoid magnets: a greatly reduced magnetic field strength for equivalent focusing. This advantage becomes even more significant in the case of beams of ion species heavier than the proton and at relativistic energies. The challenge is now to prove the reliable operation of a lens at the suitably high plasma densities and with practical high voltages and magnetic fields that are required to confine the plasma. In this context, this PhD thesis presented the identification and study of some the limitations and the capabilities of non-neutral electron plasma (Gabor) lenses.

Recent efforts at Imperial College London to produce a practical Gabor lens had limited experimental success [54] which motivated a series of simulation studies of the respective prototype lens. A central conclusion of these simulations was that the production of a stable plasma is a critical condition for the lens to produce a time-independent focusing effect, free of any aberrations. The precise ways in which unstable plasmas were recreated in particle-in-cell (PIC) simulations and the repeated observation of the anomalous focusing effect in previous experiments [54, 222] corroborate the evidence that the reduced complexity of the plasma production technique used so far can lead to deleterious effects.

Simulations of electrons confined in several plasma lenses support the suitability of the particle-in-cell method and the VSim software to make predictions on different aspects of plasma confinement (stable versus unstable time-evolution, maximum achievable plasma density, spatial distribution of the charge, plasma expansion and relaxation rates). Nonetheless, only a selection of plasma phenomena investigated in laboratory experiments were reproduced qualitatively by the simulations and the comparisons underlined the importance of further quantitative validation of the code. Moreover, the plasma phenomena studied in this work were limited to time scales

of tens of microseconds by the large computing power needed to simulate properly high-density magnetised plasmas. The VSim package and the associated code was found to be configurable to allow the adaptation of the field solver, initial particle phase-space distribution, definition of external fields, and post-processing tools. However, it was most often non trivial or ineffective to gain insight into very specific aspects of the particular implementation of the PIC method in VSim, which provides additional incentive for code validation and exploration of alternative PIC codes.

In order to track protons and ions through a plasma lens, two separate models of the space-charge field generated by the plasma were devised and tested. While the first model assumed an ideal, uniform plasma with hard edges, the second model was based on the numerical solution of equations that describe a non-neutral plasma in a state of global equilibrium. Even though the latter model resulted in a more realistic focusing field generated by the space-charge, both models produced almost identical results when used in proton-tracking simulations of LhARA. The potential benefit of using plasma lenses in other accelerator beam-lines may be an incentive to integrate one of the presented models in existing beam-tracking codes to allow more rapid investigation of plasma focusing as an alternative to the more conventional magnetic or electrostatic beam-focusing devices. It must be noted that for very intense ion beams, the effect of a plasma lens can be accounted for by a space-charge neutralisation factor in those codes that calculate space-charge effects. In the case of lower beam currents, a more elaborate model of the plasma lens is required. The models of plasma lenses devised in this work served a series of simulations of the LhARA Stage 1 beam-line. The simulations used a realistic proton distribution at the source generated via the TNSA mechanism and led to two main observations.

Firstly, the three plasma lenses at the front end of the beam-line have a capability to capture and focus the protons from the source superior to that of the equivalent solenoids. As a result, a larger number of protons are transmitted downstream when using plasma lenses compared to the case of replacing the lenses with equivalent solenoid magnets. While the above result was found in the absence of any beam collimation, the presence of the vacuum nozzle between the source and the first lens reduces the benefit of using plasma focusing over magnetic focusing. Consequently, the design of the vacuum nozzle needs to be revisited.

Secondly, effective reduction of the wide energy spread of the protons at the source was achieved in simulations with a series of beam collimators. In this case, the number of protons per bunch delivered to the end station completely within the nominal energy range was almost identical when plasma lenses or solenoids were considered. Indeed, the fraction of the protons with suitable energies that is transported from the source to the end station can be further improved by optimisation of the size and shape of the collimators. Moreover, the good capability of the beam-line to select a narrow energy interval from the wide energy spectrum produced at the source motivates the investigation of possible RF manipulation of the longitudinal phase-space of the proton bunches to achieve even higher bunch charges delivered at the end station. It must be noted that a future study should also focus on the capability of the present energy-selection scheme to filter protons or other ion species from a mixture of charged particles produced at the source. For this, PIC simulations of the source may provide realistic distributions of other ion species, such as carbon, in addition to those of protons, at the start of the

beam-line.

In addition to simulation studies, a series of preliminary measurements were performed on low-density electron clouds confined within a Penning-Malmberg trap. As the electrons were injected from an external source, the experiments emphasised the need for a source that is capable of producing large enough electron currents in order to achieve high electron densities in the trap and generate clouds of charge in a well-defined plasma regime. Overall, the destructive techniques that were used to diagnose the electron clouds were found to be a valuable tool to investigate several aspect of plasma confinement such as radial profiles, expansion rates, total particle number, or the magnetron motion. The experimental system that monitors and controls the positron/electron beam-line on which the measurements were carried out provided remarkable reproducibility of the trapped electron clouds. These techniques can be exploited in the construction of future test benches for plasma lenses and, eventually, in the control systems of plasma lenses in operation at larger facilities.

Lastly, another aspect considered in this work was the use of solenoid magnets in LhARA as an alternative to the plasma lenses. In particular, a design study was presented to evaluate the requirements of a solenoid based on a normal-conducting copper coil that can generate the strong focusing required by LhARA. It was found that such a magnet can be constructed according to analytical models with the cost of large power consumption ($>100\,\text{kW}$ per magnet), large mass (–2 t), and an intricate piping scheme for cooling. While these requirements found well-established solutions in large accelerator facilities, it remains to be determined whether normal-conducting technology is the most suitable in the case of LhARA. The large mass and power consumption of the solenoids could be reduced by using superconducting (SC) coils for which the required magnetic flux density is closer to the lower bound of the range of typical field strengths generated by SC magnets. However, the cryogenic system needed for the SC solenoids adds to the overall complexity of construction and operation of the LhARA facility. In addition, high-power lasers hitting a solid target can produce strong electromagnetic (EM) pulses. For LhARA, an investigation is needed to evaluate the effects of the EM pulses on the cryogenic system of a solenoid placed close to the particle source. Even beyond the LhARA project, the novel technologies that can contribute to the provision of proton and ion therapy on a larger scale need to evade the current limitations of high cost and complexity of operation. In this context, the plasma lenses remain the most cost-effective solution, provided that their reliable operation is proven.

In summary, the studies presented in this thesis show that plasma instabilities represent an important factor in the development and production of a Gabor lens that can serve in a research facility as a reliable device designed to focus charged particle beams. It was shown that insufficient control over the mechanism by which the plasma is produced in the lens may lead to unstable evolution of the plasma and, in turn, to deleterious focusing effects on the beam. Predictions of the dynamics of electrons confined by a plasma trap were made under several scenarios with the particle-in-cell simulation method and a dedicated software. Some aspects of the simulation method were compared against measurements of low-density confined electron clouds. In a stable regime, the plasma lenses were proven in particle-tracking simulations to provide efficient capture and

transmission of laser-driven protons at 15 MeV and to contribute significantly to a scheme which can reduce the wide energy spread of the beam within the limits suitable for radiobiology studies. For this, a couple of models were created for tracking particles through realistic space-charge fields generated by plasma lenses. Moreover, an evaluation was presented of the performance and principal requirements of solenoid magnets that provide equivalent focusing and could replace the plasma lenses in LhARA. A preliminary design for a normal-conducting solenoid has been proposed.

The next step for the development of a plasma lens that is suitable for LhARA will be the construction of a test bench and a new prototype lens. The design of the new lens should incorporate some of the techniques that are well-known within the low-energy, non-neutral plasma community and have been explored before to address the negative impact of the following factors on the stability of well-confined plasmas: long plasma length, large plasma-to-wall radii ratios, low confining magnetic field strengths, plasma densities at large fractions of the Brillouin limit. For this, an R&D programme of experiment, simulation, and design has started as part of the development of the Conceptual Design Report for LhARA. Elsewhere, efforts are being directed into increasing the scale of both Gabor-type [277] and active [39] plasma lenses, with apparatuses getting closer to some of the specifications required by a research facility based on laser-driven ions.

The promising capability of particle-in-cell simulations to provide predictions on the time evolution, spatial distribution, and maximum achievable density of plasmas produced by space-charge lenses should continue to be investigated. Current simulations still leave the door open to explore a variety of trapping challenges including the dynamic processes involved in loading a typical Penning-Malmberg trap (e.g. two-stream instability), magnetic field perturbation, resistive cooling effects, heat transfer and radial particle transport/loss. While the commercial code VSim proved to be robust and largely configurable in building simulation models and investigating the time evolution of magnetically-confined electron plasmas, some limitations have been identified. The simulations developed so far with VSim were set up to ensure that the cyclotron motion of the electrons was properly resolved, which, in turn, limited the time duration of the plasma phenomena of interest to $\sim 20\,\mu s$. Further validation of VSim in the absence of the above condition is required, in addition to the exploration of other particle-in-cell codes, to develop confidence in the results of simulations of extended equilibrium of the large, high-density plasmas required by the Gabor lens for LhARA.

Future designs for a plasma lens should also include the evaluation of strategies to generate the plasma under well controlled conditions. Given the good access to the volume where the plasma is contained in a Penning-Malmberg trap, one possible solution is to feed the electrons on-axis from an external source. This technique would require additional bending magnets to guide the electron beam into the lens and studies are required regarding the integration of these magnets with the main beam-line for the positively charged protons or ions. Furthermore, the time it would take for the nominal plasma density to be established each time the lens is filled needs to be evaluated. The loading time, together with the duration for which the plasma remains stable and the plasma density does not start to decrease, could place constraints on the repetition rate and duty cycle

with which the proton and ion bunches can be delivered to the end stations in LhARA. It must be noted that well established techniques exist to achieve steady-state plasma densities and long confinement times (for example, the Rotating Wall, RW [212, 215, 240]), and a combination of these methods led to unprecedented control of trapped positrons and electrons [278].

Another important aspect regarding the use of electron plasma lenses for focusing proton and ion-beams is the question of how close to the Brillouin (the maximum possible) density limit the plasma can be maintained in a steady state by a future lens. With the RW technique, the Brillouin limit has been approached within a factor of six or less [216] for comparatively low trap magnetic fields (< 0.1 T). At higher magnetic fields, higher absolute plasma densities have been achieved in high-field (> 1 T) UHV traps, but the corresponding fractions of the Brillouin limit achieved were much smaller [216] (typically only $<1\%$). While the relatively poor performance of the traps at present in the high-field regime is a subject of ongoing research, the identification of the limiting factors play an important role in evaluating the capability of plasma lenses to focus proton and ion beams with kinetic energy above ~ 100 MeV. One of the leading advantages of the plasma lens is that the magnetic field strength required to confine a plasma with a certain density is smaller than that of a solenoid that would give equivalent focusing by a factor in the range 10–100 if one assumes that plasma densities close to the Brillouin limit can be achieved. However, if only a fraction, χ_B, of the density limit can be achieved, then the lens would require a magnetic field larger by a factor of $\sqrt{\chi_B}$ and the advantage of space-charge focusing over focusing by solenoid magnets would become less significant.

In conclusion, a shift of paradigm has unfolded in the last couple of years from the efforts to develop a plasma lens for LhARA and beyond. These efforts turned the 'single-knob-lens' concept that has been experimented upon for many decades into a systematic plan to integrate and test techniques, which have been proven separately. Due to this plan of experiment- and simulation-guided design decisions, there is reason to expect a step change in the reliability and capabilities of non-neutral plasma lenses.

Appendix A

Other historical developments of plasma lenses

Plasma lenses for focusing heavy ions

In parallel to the development programme in Germany, substantial experimental investigations have also been reported from the Institute of Physics NAS of Ukraine starting from around 1990 [279]. The goal was to develop a practical plasma lens dedicated to the production of high-current (up to a few Amperes), low-divergence heavy-ion beams. The work followed the initial confirmation of the focusing capability of a plasma lens at low ion-beam currents more than twenty years earlier [280]. In the first reports, it was shown that the magnetic field geometry, the external potential distribution applied to the lens electrodes and the ion beam current have an impact on the radial profile of the total electric potential [281]. Thus, it was possible to tailor the radial focusing force to reduce the aberrations caused by the lens.

With the optimal configuration of the lens found earlier with hydrogen ion beams, the effectiveness of the plasma lens was further demonstrated for copper ions [282]. It must be noted that the magnetic field of the lens was operated in a pulsed mode at the same frequency as that of the source (1 Hz). The background gas pressure was chosen such that the plasma creation could be done only by secondary emission of electrons caused by the collision of the beam ions with the lens electrodes. Discharge electron currents of up to 200 A were measured on the lens electrodes during the passage of the beam and the duration of the discharge was found to be greater than the duration of the beam pulse. It is also interesting to note that the lens working residual-gas pressure was an order of magnitude higher for copper ions passing through than in the case of hydrogen beams. Although the lens was successfully shown to be capable of increasing the beam-current densities delivered to a target by up to about an order of magnitude, the focusing strength of the lens was strongly dependent on the ion-beam current passing through it. As the focusing capabilities of the lens were confirmed with other ions as well (Zn, Mo, C, Ta) [45, 283, 284], the use of the plasma was proposed early on as part of a high-dose ion-implantation facility.

The working principle of these 'electrostatic plasma lenses' (EPL) developed jointly by the Institute of Physics in Kyiv and LBNL in Berkley is based on a generalisation of the initial space-charge lens proposed by Gabor. The design of an EPL is similar to that of the Gabor lens: a series of cylindrical ring electrodes are located within an externally driven magentic field, with field lines connecting ring electrode pairs symmetrically about

the lens mid-plane [285]. Electrons produced inside the volume of the lens are able to stream freely along the field lines, thereby tying the potential to that of the electrostatic ring to which the magnetic field line is attached. Thus, the magnetic field lines are also equipotential lines or, equivalently, the electric field is perpendicular to the magnetic field. The resulting focusing force is electrostatic in nature since the focal length of the lens does not depend on the charge-to-mass ratio of the ions in the beam [285]. If the beam current density is low, the electron plasma in the lens acts to neutralise the beam space-charge and to compensate for the vacuum potential of the lens to make the electric field lines transverse to the magnetic field lines. For a high current beam, the electrons of lower density only partially neutralise the space-charge of the beam. In this case, the electrons are held within the electric potential of the beam. Nonetheless, the motion of the magnetised electrons leads to the equpotentialisation described above. Lens operation in the latter regime occurs when [285]

$$\frac{I_b}{4\pi\epsilon_0 v_b} \gg V_{\mathrm{el,max}}, \tag{A.1}$$

where I_b, v_b are the ion beam current and velocity, and $V_{\mathrm{el,max}}$ is the maximum voltage applied to one of the lens electrodes. The condition of equipotentials along the magnetic field lines implies that the distribution of the magnetic field is critical to obtain a radial electric field that is linear and does not create focusing aberrations.

A second lens was designed and built by the same research group by replacing the typical electrically-driven solenoid of the lens with permanent magnets which have a lower cost [286]. The configuration of the magnets was calculated with computer simulations to produce the required magnetic field shape. Variations in the shape of the field could be made by using iron pieces to shunt parts of the field. It was experimentally confirmed that the focusing properties of the lens are determined by the shape of the magnetic field and the distribution of the external potential applied to the ring electrodes. The optimal configurations were found empirically on a case-by-case basis to minimise the aberrations. It was also reported that the focusing properties of the lens were completely lost or showed significant variations from pulse to pulse for strengths of the magnetic field below a certain threshold around $0.01\,\mathrm{mT}$. However, a different report [287] describes a significant improvement in the focusing properties of the lens for a very narrow range of low magnetic field. In this range, the beams experienced a stronger focusing effect and the collective oscillations measured within the lens volume were strongly suppressed. With increasing magnetic field, this regime is quickly destroyed. Numerical simulations suggested that for a magentic field higher than optimal, the finite width of the lens electrodes leads to the formation of layered electron structures along the field lines [288] which limit the focusing strength.

Finally, measurements of the transport of copper ions through the permanent-magnet plasma lens showed that the beam emittance and noise do not increase for optimal lens settings [289, 290]. Since the formation of the electron plasma in an EPL is induced by the ion beam pulse, the maximum focusing effect is observed near the centre of the beam pulse. This modulation of the focusing strength during the duration of the beam pulse might be undesirable for possible applications of the lens. Given the extent of the investigations of several electrostatic plasma lenses for heavy-ion beam focusing, a recent review of the development is provided in [291].

Plasma lenses for focusing negative ion beams

In parallel to the experimental efforts, space-charge lenses were also theoretically compared with electrostatic quadrupoles [226]. It was also realised at that time that positron plasmas that could effectively focus negative ion beams would be difficult to achieve due to the many orders of magnitude difference between the required positron densities and the experimentally achievable ones. While much higher positron densities can be confined today [187], the large total number of positrons required to achieve long plasma columns or centimetre-level plasma radii would require long trap filling times with the current moderated positrons obtained from radioactive sources.

In the redesign of one of the Fermilab linacs, a non-neutral plasma lens was also considered for matching an H^- beam into an RFQ [292]. From the theoretical considerations, it was found that magnetic fields of at least a couple of Tesla were required for the confinement of a non-neutral proton or positive ion plasma that could have focused the H^- beam available from the magnetron source in use at that time. A magnet that could create such a large field was considered to be difficult to construct, would have used a large amount of power, and would have produced a focusing effect at least as strong as the Gabor lens. As a result, focusing negative ion beams with space-charge lenses was considered not to be feasible. As part of the numerical considerations in [292], it was already established at that time that the typical densities of plasmas that were confined were a factor of five or more below the Brillouin limit.

As an alternative to a positron plasma, the use of confined positive ions to focus $\sim 10\,\mathrm{keV}$ H^- beams was proposed and experimentally investigated by [293]. The working-principle of the lens was based on the H^- beam ionising a gas inside the volume of the lens to create a mixture of positive ions and free electrons. Two transparent electrodes at the ends of the lens produced an electric field that extracted the electrons out of the lens [294]. Due to their larger inertia, the positive ions would remain inside the lens to form the space-charge required to focus the negative beam. During the experiments, Ar, Kr, Xe were chosen as working gasses because of their relatively high ionisation cross-sections. As gas has to be continuously supplied to the lens, there is a lower gas pressure limit to ensure focusing which is achieved when the space charge of the positive ions formed by ionisation exceeds the space charge of the negative beam. Under this condition, the focusing strength was measured to increase with the background gas pressure [293]. However, the required gas pressures were high enough that the loss of beam ions from collisions with neutrals was critical. The lens tested in [293] was modified to work at lower pressure by replacing the production of the space-charge by the beam with a non-self-maintained plasma discharge. To achieve that, the lens diameter was increased such that electrons formed at the surface of the electrodes could be accelerated over a large-enough distance up to an energy sufficient for gas ionisation. Further improvement in the focusing strength of the lens was achieved by placing an incandescent tungsten wire inside the lens to act as an electron emitter and an additional driver of gas ionisation [295].

Measurement of plasma density

With the application of RF plasma diagnostic techniques, it was soon possible to investigate the changes in the confined-plasma density induced by variations in the confining fields [52]. The general trends are applicable to any plasma confined in a Penning-Malmberg type trap. The electron density increases with both the external voltage and the strength of the magnetic field (B) of the trap when the trap voltage (V_a) is high compared to the self-electrostatic potential of the plasma. If the magnetic field is increased to high values, the electron density reaches saturation. Under a relatively low trap voltage, the plasma density reaches saturation more easily at low magnetic fields [209]. In other words, for $V_a \ll \Delta\phi$ (see eq. (3.5)) the electron density increases with V_a and B, while for $V_a \approx \Delta\phi$, the electron density only depends on V_a.

In some of the earliest space-charge lenses, the electron density was measured using an RF drive frequency to excite waves in the plasma and subsequently measuring the plasma response [206, 209, 296]. An azimuthally-symmetric longitudinal plasma wave is typically excited by driving an RF voltage on a ring electrode close to one of the ends of the plasma and the response signal is picked up with a similar ring electrode at the opposite end of the electron column. Then, a spectrum analyser can be used to extract any relevant resonance frequencies. Sometimes, it is possible to pick up a sharp signal in the absence of a drive voltage [206]. In principle, this signal could mark the rotation frequency of the plasma as a rigid column around the symmetry axis. To first order, the rotation frequency is independent of the column length and radius and can act as a direct measurement of the plasma frequency (see eq. (3.9)). If a driving signal is coupled into the trap, the resonant response picked-up from the plasma corresponds to the lowest modes of electrostatic standing waves in the plasma [297]. The measured frequencies can then be used to determine the electron density [187, 209]. Tapered ring electrodes [206, 296] were employed to ensure that the surface of the plasma follows the shape of a magnetic-flux tube inside the trap and that the excited waves follow a simpler dispersion relation that could result in more accurate determination of the plasma density.

A second method that was used to measure the electron density relies on the positive residual gas ions that are expelled from the lens. The trapped electrons reduce the on-axis potential while the residual gas ions produced are longitudinally accelerated out of the lens. By measuring the energy of the emitted ions, one can determine the average electron density [53]. Multiple distinct ion energies can be measured due to the residual gas ions being produced in different regions of the lens, for example, near the anode or near the symmetry axis.

Appendix B

Single-particle linear dynamics for a selection of beam-focusing lenses

B.1   Single-particle Hamiltonian and focusing strength

The equations of motion for a charged particle moving through a focusing field allows the dependence of the focusing strength on the key parameters of the beam-transport system to be extracted. Furthermore, a comparison can be made of different focusing systems for beams with specific ion species and momentum or energy. The approach chosen here follows the treatment of [298] where linear equations of motion are obtained from an approximation to the Hamiltonian which is expanded as a power series to second order in the dynamical variables.

The dynamics of a particle is in general defined by the Hamiltonian, H, which is a function of the co-ordinates, x_i, of the particle, the components, P_i, of its momentum, and time t, $H = H(x_i, P_i; t)$. The equations of motion of the particle are produced from Hamilton's equations

$$\frac{dx_i}{dt} = \frac{\partial H}{\partial P_i}, \qquad \frac{dP_i}{dt} = -\frac{\partial H}{\partial x_i}. \tag{B.1}$$

In an accelerator beam-line, we are interested in charged particles moving through electromagnetic fields. The electric and magnetic fields, $\mathbf{E}$ and $\mathbf{B}$, can be expressed with the use of a scalar potential ϕ and vector potential $\mathbf{A}$ from the euqations

$$\mathbf{E} = -\nabla\phi - \frac{\partial \mathbf{A}}{\partial t}, \qquad \mathbf{B} = \nabla \times \mathbf{A}. \tag{B.2}$$

An appropriate Hamiltonian that gives the correct equations of motion consistent with the Lorentz force for a relativistic particle with charge q and mass m is [298]

$$H_1 = -\sqrt{\frac{(E - q\phi)^2}{c^2} - (p_x - qA_x)^2 - (p_y - qA_y)^2 - m^2c^2} - qA_z. \tag{B.3}$$

Note that with the Hamiltonian (B.3), p_i are the components of the relativistic canonical momentum $\mathbf{p} =$

$\beta\gamma mc + q\mathbf{A}$. Using the Hamiltonian (B.3), the equations (B.1), and knowing the electromagnetic potentials as functions of position, one can write down the equations of motion for a charged particle moving along a straight accelerator beam line.

However, in practice, the motion of the particles is described with respect to a reference particle with a fixed reference momentum, $P_0 = \beta_0\gamma_0 mc$, for which the beam line has been designed. Here, β_0, γ_0 are the relativistic factors. A new Hamiltonian is obtained after a sequence of canonical transformations [298] of the dynamical variables to ensure that Hamilton's equations still describe the evolution of the system. The new Hamiltonian is

$$H = \frac{\delta}{\beta_0} - \sqrt{\left(\delta + \frac{1}{\beta_0} - \frac{q\phi}{cP_0}\right)^2 - (p_x - a_x)^2 - (p_y - a_y)^2 - \frac{1}{\beta_0^2\gamma_0^2}} - a_z. \tag{B.4}$$

As part of the canonical transformations, the Hamiltonian is scaled by the reference momentum, P_0, which requires an equivalent scaling of the canonical momentum and the vector potential as

$$p_{x,y} = \frac{\beta_{x,y}\gamma mc + qA_{x,y}}{P_0}, \tag{B.5}$$

$$\mathbf{a} = \frac{q}{P_0}\mathbf{A}. \tag{B.6}$$

Finally, a new longitudinal variable z is defined as $z = s/\beta_0 - ct$ from the independent variable s which is the distance along the beam-line. In this way, the coordinate z is approximately the longitudinal distance between the particle and the reference particle. This transformation of the longitudinal coordinate requires an equivalent modification of the longitudinal momentum which becomes the energy deviation

$$\delta = \frac{E}{cP_0} - \frac{1}{\beta_0}. \tag{B.7}$$

In reality, the equations of motion of particles in all accelerator components are nonlinear. However, many important aspects of beam dynamics can be understood in terms of the properties of these components using a linear approximation. Expanding the Hamiltonian (B.4) as a power series to second order in the dynamical variables and substituting it into Hamilton's equations (B.1) produces linear equations of motion. The truncation of a power series expansion of the Hamiltonian at low order in p_x and p_y is known as the paraxial approximation. If the equations of motion can be solved exactly under this approximation, the solutions can be expressed in terms of a transfer map, R, such that

$$\vec{x}_1 = R\vec{x}_0, \tag{B.8}$$

where $\vec{x}_0, \vec{x}_1$ are the phase space vectors $\vec{x} = (x, p_x, y, p_y, z, \delta)^\mathsf{T}$ at the entrance and exit of the beam line component, respectively. The rest of this subsection evaluates the transfer maps for some of the accelerator components discussed previously that could be use to capture and focus laser-driven proton and ion beams.

Solenoid

A charged particle moving through a solenoid magnet experiences a uniform magnetic field that is parallel to the direction in which the beam is travelling. The magnetic field produced by a solenoid is given by $\mathbf{B} = (0, 0, B_0)$. A conventional choice of the equivalent vector potential is

$$\mathbf{A} = \left(-\frac{B_0}{2} y, \frac{B_0}{2} x, 0 \right). \tag{B.9}$$

Special care is needed at the entrance and exit faces of the solenoid where the dynamical variables are transformed from a field-free region to the region within the magnet. If we assume that the magnetic field changes within a fringe field region that is infinitesimally short, the corresponding transfer map for the fringe fields is the identity [298]. With the vector potential (B.9), the Hamiltonian (B.4) for a solenoid is

$$H = \frac{\delta}{\beta_0} - \sqrt{\left(\delta + \frac{1}{\beta_0} \right)^2 - (p_x + k_s y)^2 - (p_y - k_s x)^2 - \frac{1}{\beta_0^2 \gamma_0^2}}. \tag{B.10}$$

where the focusing strength k_s is defined by $k_s = (qB_0)/(2P_0)$. Expanding the Hamiltonian to second order in the dynamical variables gives

$$H_2 = \frac{p_x^2}{2} + \frac{p_y^2}{2} + \frac{\delta}{2\beta_0^2 \gamma_0^2} + \frac{k_s}{2} x^2 + \frac{k_s}{2} y^2 - k_s x p_y + k_s y p_x. \tag{B.11}$$

The fourth and fifth terms in (B.11) act as focusing terms which shows that the solenoids provides focusing simultaneously in both transverse directions. Furthermore, the focusing is independent of the sign of k_s, so independent of the direction of the field as parallel or antiparallel to the direction of propagation of the beam. The final two terms in (B.11) show that there is a coupling between the horizontal and vertical motions which is a result of the particles moving in helical trajectories by performing cyclotron rotations around the field lines.

Solving Hamilton's equations given the second-order Hamiltonian B.11 results in the transfer matrix for a solenoid with length L:

$$R_{sol} = \begin{pmatrix} \cos^2(\omega L) & \frac{1}{2\omega}\sin(2\omega L) & \frac{1}{2}\sin(2\omega L) & \frac{1}{\omega}\sin^2(\omega L) & 0 & 0 \\ -\frac{\omega}{2}\sin(2\omega L) & \cos^2(\omega L) & -k_s\sin^2(\omega L) & \frac{1}{2}\sin(2\omega L) & 0 & 0 \\ -\frac{1}{2}\sin(2\omega L) & -\frac{1}{\omega}\sin^2(\omega L) & \cos^2(\omega L) & \frac{1}{2\omega}\sin(2\omega L) & 0 & 0 \\ \omega\sin^2(\omega L) & -\frac{1}{2}\sin(2\omega L) & -\frac{\omega}{2}\sin(2\omega L) & \cos^2(\omega L) & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \frac{L}{\beta_0^2\gamma_0^2} \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \tag{B.12}$$

where $\omega = \sqrt{k_s}$ and the focusing strength parameter k_s is

$$k_s = \left(\frac{q}{P_0} \frac{B_0}{2} \right)^2. \tag{B.13}$$

It is worth noting that making a second-order approximation to the Hamiltonian for a particle in a solenoid field, the focusing strength experienced by a particle is independent of the energy deviation δ. The variation of the focusing strength with particle energy in an accelerator beam line is known as chromaticity. The chromaticity is a nonlinear effect and is properly described by including higher-order terms in the transfer map. When designing a focusing element for laser-generated ion beams, chromaticity may play a role in the design of the energy selection scheme. The change in the focusing strength with the particle energy can be employed in combination with properly placed collimators to select a specific energy interval from the continuous energy spectrum produced at the source. The chromatic dependence on the beam energy can be recovered from eq. (B.13) by choosing different reference momenta and zero energy deviation for each particle with a different momentum.

Non-neutral electron plasma lens

In a non-neutral electron plasma lens of the Gabor type [49], the focusing force is given by the radially inward electrostatic attraction of the negative space-charge of the electrons. An ideal linear focusing force is obtained for a spatially-uniform plasma density n_e. Such a cylindrically symmetric uniform electron plasma produces a radial electric field

$$E_r = -\frac{en_e}{2\epsilon_0}r, \tag{B.14}$$

where r is the distance form the symmetry axis. The electric field, E_r, is equivalent to an electrostatic potential, $\phi = (en_e r^2)/(4\epsilon_0)$, if we assume an infinitely long plasma column. In a plasma lens with a finite length, the particles are accelerated and decelerated such that there is no net change in energy after passing through the lens. Plugging the potential ϕ into eq. (B.4) and neglecting any edge effects, the Hamiltonian for a non-neutral electron plasma lens is

$$H = \frac{\delta}{\beta_0} - \sqrt{\left[\delta + \frac{1}{\beta_0} - \frac{qen_e}{4\epsilon_0 cP_0}(x^2 + y^2)\right]^2 - p_x^2 - p_y^2 - \frac{1}{\beta_0^2\gamma_0^2}}. \tag{B.15}$$

In order to obtain a linear approximation to the solutions to the equations of motion, the space-charge potential of the plasma is considered to be much smaller than the reference beam energy normalised by the charge, i.e. $q\phi \ll cP_0$, in addition to the usual paraxial approximation, $p_x, p_y \ll 1$. In practice, the former assumption is correct since metre-scale focal lengths can be achieved from a non-neutral electron plasma by using longitudinally-confining voltages of several tens of kilo-Volts. Expanding the Hamiltonian to second order in the dynamical variables and dropping the constant terms gives

$$H_2 = \frac{p_x^2}{2} + \frac{p_y^2}{2} + \frac{\delta^2}{2\beta_0^2\gamma_0^2} + \frac{k}{2}(x^2 + y^2), \tag{B.16}$$

where the focusing parameter, k, is defined by

$$k = \frac{qe}{2\epsilon_0}\frac{m\gamma_0}{P_0^2}n_e. \tag{B.17}$$

Solving the equations of motion that follow from the Hamiltonian (B.16) gives the transfer matrix R_{GPL} for a distance L through the Gabor-type plasma lens:

$$
R_{GPL} = \begin{pmatrix}
\cos(\omega L) & \frac{\sin(\omega L)}{\omega} & 0 & 0 & 0 & 0 \\
-\omega \sin(\omega L) & \cos(\omega L) & 0 & 0 & 0 & 0 \\
0 & 0 & \cos(\omega L) & \frac{\sin(\omega L)}{\omega} & 0 & 0 \\
0 & 0 & -\omega \sin(\omega L) & \cos(\omega L) & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & \frac{L}{\beta_0 \gamma_0} \\
0 & 0 & 0 & 0 & 0 & 1
\end{pmatrix},
\tag{B.18}
$$

where $\omega = \sqrt{k}$.

Active plasma lens

The active plasma lens produces an azimuthal magnetic field by passing a strong axial sub-μs current pulse through a gas-filled capillary. After the gas is fully ionised, the plasma reaches a quasi-equilibrium characterised by a uniform current density. The radial magnetic field gradient, $\partial B_\theta / \partial r$, within the aperture is then constant and it is given by eq. (1.5). The resulting magnetic field inside the plasma channel is related to the gradient, g, by

$$
\mathbf{B} = (-gy, gx, 0),
\tag{B.19}
$$

for the current propagating in the same direction as the beam (positive z direction). The magnetic field can be obtained from the vector potential

$$
\mathbf{A} = \left(0, 0, -\frac{g}{2}(x^2 + y^2)\right).
\tag{B.20}
$$

Substituting into the general accelerator Hamiltonian B.4, the particle motion inside the APL is obtained from

$$
H = \frac{\delta}{\beta_0} - \sqrt{\left(\delta + \frac{1}{\beta_0}\right)^2 - p_x^2 - p_y^2 - \frac{1}{\beta_0^2 \gamma_0^2}} + \frac{gq}{2P_0}(x^2 + y^2).
\tag{B.21}
$$

To obtain the transfer matrix for the APL, H is expanded to second order in the dynamical variables and the zeroth-order term is dropped to obtain

$$
H_2 = \frac{p_x^2}{2} + \frac{p_y^2}{2} + \frac{\delta^2}{2\beta_0^2 \gamma_0^2} + \frac{k}{2}(x^2 + y^2),
\tag{B.22}
$$

where the focusing parameter depends on the field gradient as

$$
k = \frac{\partial B_\theta}{\partial r} \frac{q}{P_0}.
\tag{B.23}
$$

The transfer matrix R_{APL} has an identical expression to (B.18) with the focusing parameter k suitably replaced.

Thin-lens approximation

The analysis of the focusing capabilities of the beam-line components described above can be simplified when the resulting focal length f is much larger than the length of the focusing field L. Under this condition, the 'thin lens' approximations is employed by taking the limit $L \to 0$, while the product (kL) remains fixed such that

$$kL \to \frac{1}{f}, \tag{B.24}$$

where k is the focusing strength parameter of any of the lenses considered in this section. For capturing ion beams with energy $>10\,\mathrm{MeV}$, metre-scale focal lengths are needed to keep the beam-transport line compact. In this case, the focal lengths are comparable to the length of a solenoid or a non-neutral plasma lens required to provide such strong focusing, and the thin lens approximation is not appropriate anymore. A comparison between different focusing techniques can be achieved based on how the focusing strength parameter k of each device depends on its design characteristics or, alternatively, with the use of beam-tracking simulations.

B.2 Single-particle motion in a solenoidal magnetic field

The typical textbook treatment of how charged particle beams are focused by a solenoid magnet employs the evolution of the beam envelope [127], a paraxial approximation of the single-particle Hamiltonian [298] or the application of Busch's theorem [299] of the conservation of canonical angular momentum in the axisymmetric magnetic field of the solenoid. A more graphically descriptive approach by Kumar [300] is adapted and presented in the following to emphasise the focusing effect of a solenoid on the motion of a single charged particle, as summarised in fig. B.1.

In the absence of sources of charge and current, electric and magnetic fields may be derived from a scalar potential that obeys the Laplace equation. In an axisymmetric system, the full potential distribution may be determined from the potential distribution and its derivatives on the z-axis ($r = 0$) [127], where r is the radial distance from the solenoid axis. Since the magnetic potential is not a measurable quantity, the field components B_r and B_z are rather obtained from the axial field component $B_z(0, z) = B(z)$ on the axis, where z is the distance along the solenoid axis. The series representation of the field components gives

$$B_z(r, z) = \sum_{n=0}^{\infty} \frac{(-1)^n}{(n!)^2} \frac{\partial^{2n} B}{\partial z^{2n}} \left(\frac{r}{2}\right)^{2n} = B(z) - \frac{r^2}{4} \frac{\partial^2 B(z)}{\partial z^2} + \cdots \tag{B.25}$$

and

$$B_r(r, z) = \sum_{n=1}^{\infty} \frac{(-1)^n}{n!(n-1)!} \frac{\partial^{2n-1} B}{\partial z^{2n-1}} \left(\frac{r}{2}\right)^{2n-1} = -\frac{r}{2} \frac{\partial B(z)}{\partial z} + \frac{r^3}{16} \frac{\partial^3 B(z)}{\partial z^3} + \cdots . \tag{B.26}$$

For simplicity, we assume a cylindrical beam of protons with a uniform distribution of particles. The protons have rest mass m_p, charge e and move initially parallel to the beam axis with velocity $v = v_z$. In the paraxial approximation, we consider the beam size to be small in comparison to the diameter of the solenoid and linearize

the field components. The first-order magnetic field terms are

$$B_r = -\frac{1}{2}\frac{\partial B(z)}{\partial z}, \quad B_z = B(z). \tag{B.27}$$

A further simplification is to consider that the solenoid of length L_s produces a uniform magnetic field for $0 < z < L_s$ that abruptly drops to zero at the edges of the magnet and has no radial component. From eq. B.27 the expression for the magnetic field in this case can be written as

$$B_z(z) = B_0[H(z) - H(z - L_s)], \tag{B.28}$$

$$B_r(z) = -\frac{r}{2}B_0\left[\delta(z) - \delta(z - L_s)\right], \tag{B.29}$$

where $H(z)$ is the Heaviside step function, i.e. $H(z) = 1$ for $z > 0$ and $H(z) = 0$ otherwise, and $\delta(z)$ is the Dirac delta function which approximates the fringe fields of the solenoid. When a proton enters the solenoid, it passes the boundary between the field-free region and the uniform-field region and it experiences an azimuthal kick due to the Lorentz force. Due to this impulse, the azimuthal velocity v_θ is increased by

$$\Delta v_\theta = r_0\frac{eB_0}{2\gamma m_p} = r_0\frac{\omega_c}{2} = r_0\omega_L, \tag{B.30}$$

where r_0 is the radial coordinate of the proton when it enters the solenoid, γ is the relativistic Lorentz factor, and ω_c, ω_L are the cyclotron and Larmour frequencies, respectively. Due to conservation of kinetic energy, the longitudinal component of the velocity of the proton will also change. However, in the paraxial approximation $v_\perp \ll v_z$, where $v_\perp$ is the transverse component of the particle velocity. Hence, we consider the change in v_z to be negligible here.

Once inside the solenoid, the proton will move along a helical trajectory with a radius equal to the cyclotron radius

$$r_c = \frac{\gamma m_p v_\theta}{eB_0} = \frac{r_0}{2}. \tag{B.31}$$

Thus, every particle that enters the solenoid off-axis will move along a curved trajectory that has a radius equal to half of the initial radial displacement of the particle from the axis of the solenoid and that intersects the central axis. Figure B.1 shows the trajectory in the transverse plane of two charged particles that propagate through a constant solenoidal field. Each particle sweeps an angle 2θ with respect to the axis of its helical path marked by O_B and an angle θ measured at the solenoid axis O. Thus, the particles rotate with the cyclotron frequency $\omega_c = eB_0/\gamma m_p$ about an axis passing through the centre of their individual circular trajectories, but circle the solenoid axis with the Larmour frequency $\omega_L = \omega_c/2$. Inside the uniform field region, the radial and azimuthal
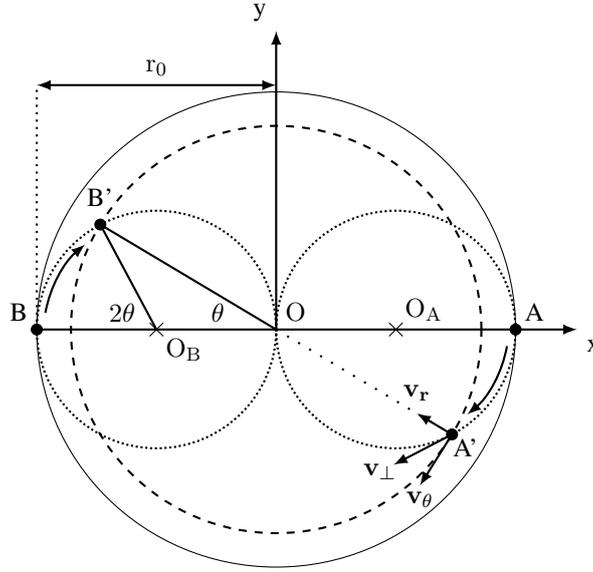
Figure B.1: Schematic representation of the focusing of a proton beam by a solenoid (B-field out of the plane of the drawing), adapted from [300] and [301]. The solid curve shows the beam envelope at the entrance of the solenoid. The dashed curve shows the beam envelope after the protons travel some distance through the solenoid. The trajectories of two test protons that enter the field at A and B are indicated by the dotted circles. The particle at B sweeps an angle 2θ with respect to the axis of its helical path marked by O_B and an angle θ measured at the solenoid axis O. Therefore, the Larmour frequency is half the cyclotron frequency. For proton A, the transverse velocity $v_\perp$ is decomposed into the radial (v_r) and azimuthal (v_θ) components.

coordinate sof the particle are given by

$$r = r_0 \cos\left(\frac{\omega_L z}{v_z}\right), \tag{B.32}$$

$$\theta = \theta_0 + \frac{\omega_L z}{v_z}, \tag{B.33}$$

where ω_0 is the initial value of θ, and the coordinates (r, θ) are measured relative to the axis of the solenoid. According to the explanation above, the protons undergo periodic focusing in the uniform field region of the solenoid. However, in the case of the applications of interest here, such as the focusing of protons with kinetic energy of several MeV, the limited length and field strength of the solenoid lead to the focal spot being located downstream of the magnet. Thus, the particles pass the fringe field at the exit plane of the solenoid as well. Equations B.32 and B.33 can be used to derive the radial and azimuthal components of the particle velocity at an arbitrary location inside the solenoid as

$$v_r = -r\omega_L \tan\left(\frac{\omega_L z}{v_z}\right), \tag{B.34}$$

$$v_\theta = r\omega_L. \tag{B.35}$$

The final stage of the focusing effect takes place when the particles interact again with a non-zero B_r when they exit the solenoid. In general, protons exit the uniform field region with a different radial displacement $r(z = L_s) = r_1 \neq r_0$. According to eq. B.35, the azimuthal velocity of the proton just before the field

boundary is $v_\theta = r_1\omega_L$. The coupling between B_r and v_z gives an additional impulse in the azimuthal direction $\Delta v_\theta = -r_1\omega_L$ similar to the initial kick described by eq. B.30. Hence, we find the important result that at the exit of the solenoid the particles emerge with no azimuthal component of velocity, i.e. $v_\theta = 0$. The radial velocity remains unchanged at the boundary and is given by

$$v_r = -r_1\omega_L \tan\left(\frac{\omega_L L_s}{v_z}\right) = -kr_1. \tag{B.36}$$

The particles move in straight lines downstream of the solenoid and gain a radial velocity towards the solenoid axis which is proportional to their radial displacement. Thus, the beam is focused to a point after travelling a distance $z = v_z/k$ from the exit plane of the magnet.

One special case corresponds to the thin lens approximation when $L_s \ll v_z/\omega_L$. In this case, $r_1 \approx r_0$ and the solenoid imparts a radial velocity to every particle given by

$$v_r = -r_0 \frac{e^2}{4\gamma^2 m_p^2 v_z} B_0^2 L_s. \tag{B.37}$$

Since the trajectory of a particle downstream of the solenoid is a straight line, the resulting equation $f_s = v_z(r_0/v_r)$ can be used to obtain the focal length f_s of the thin solenoid lens as

$$\frac{1}{f_s} = \frac{e^2}{4\gamma^2 m_p^2 v_z^2} B_0^2 L_s. \tag{B.38}$$

For an arbitrary variation of the solenoid field $B_z(z) = B(z)$ along the axis, eq. B.38 can be generalised [127] by introducing an integration of $B(z)$ over the entire length of the solenoid to obtain

$$\frac{1}{f_s} = \frac{e^2}{4\gamma^2 m_p^2 v_z^2} \int B^2 dz = \frac{e^2}{4p^2} \int B^2 dz. \tag{B.39}$$

It is important to note the dependency of the focal length on the momentum, p, of the incoming particles which shows that the beam manifests chromatic aberrations downstream of a solenoid. In the case of an incoming beam with a broad energy spectrum as the one generated by TNSA, the chromatic dependence on momentum ($p \approx \sqrt{E}$) can be used to perform a rough energy selection. In such a scheme, a collimator placed at the position where the nominal energy is brought into focus removes the particles with a large energy deviation.

The approach presented here to describe the focusing effect of a solenoid is based on assumptions that may be partially or fully violated in a real situation where the solenoid is used to capture particles from a laser-driven source. The thin-lens approximation is not valid for low energy particles and for beams with a size comparable to the aperture of the solenoid. To obtain high capture efficiency, a significant part of the initial proton flux from the source passes in close vicinity to the windings of the solenoid at large displacements from the central axis. Furthermore, the imperfections in the winding of the solenoid lead to a non-ideal current distribution that introduces perturbations in the magnetic field compared to the idealised case of eqs. B.25 and B.26.

Appendix C

Electron-trapping sequences

This appendix provides a more detailed description of the experimental sequences employed to carry out the measurements presented in chapter 5. Firstly, two procedures that were used to capture or accumulate electrons from the source into the storage trap are described. Then, an outline is given of the control system that allowed the experimental cycles to be configured easily. As the properties of the trapped electrons were measured destructively, the appendix finishes with a description of how the raw data was recorded and post-processed to extract the parameters of interest including the number of electrons and the transverse position and radial profile of the plasma.

C.1 Electron capture

Two distinct techniques were employed to capture electrons inside the volume of the particle trap from the beam emitted by the source. Figure C.1 shows two examples of camera images of the electron clouds that were obtained with each capturing technique. The precise steps of each method are described in the following.

Dynamic particle trapping

The first technique to capture electrons inside the trap is based on a fast, dynamic change of the confining potential of the trap as shown schematically in fig. C.1. Before electrons are trapped, all the electrodes of the trap are set to ground voltage. Thus, the electron beam propagates axially from the source to the MCP/P-screen assembly in front of which a negatively-biased grid blocks the continuous beam from hitting the detector. At the beginning of the capturing process, the confining potential of the trap is rapidly turned on by raising the voltage on two electrodes used as end caps. As a result, the part of the beam between the two end electrodes becomes trapped. The voltage on the blocking electrodes must be raised to a value higher than the parallel energy of the electron beam. Typically, the blocking potential was raised to the maximum available voltage of $-140\,\mathrm{V}$. It was observed that the more rapidly the potential was raised to capture the beam within the trap, the more well-defined the radial profile of the resulting electron cloud was. The time it took for the trap potential to be switched on was limited by the minimum time step required for the analogue output to change the voltage
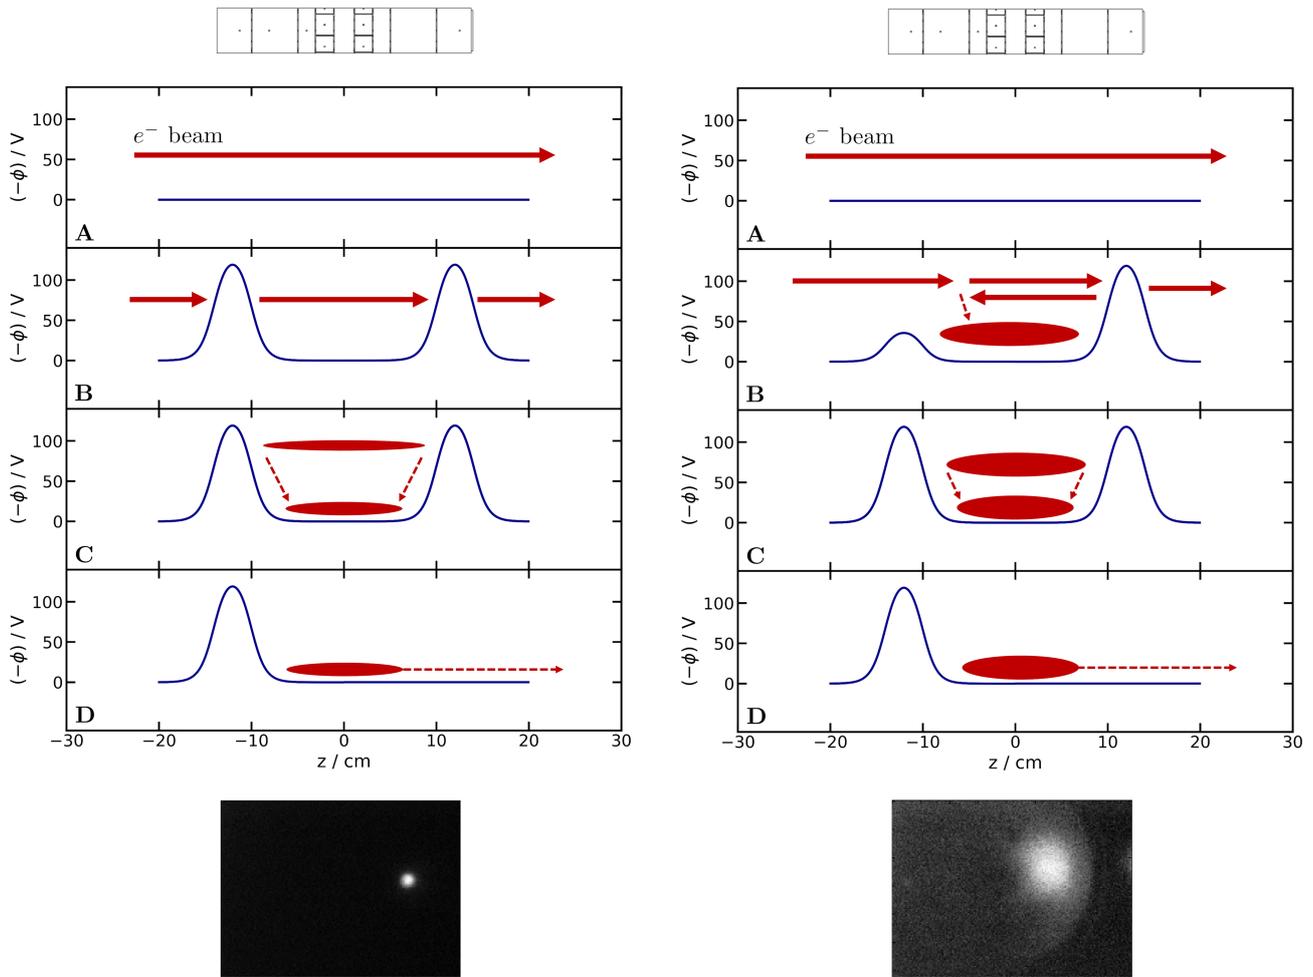
Figure C.1: Diagrammatic representation of a typical experimental sequence: A–reset of the trap potential, B–electron capture or accumulation by raising the voltage on the two end-cap electrodes to a maximum of $-140\,\mathrm{V}$, C–electron cloud storage and study, D–electron cloud release by lowering the voltage on the end of the trap closer to the electron collector. The electron streams are indicated by the red arrows. The shape of the axial electrostatic potential of the trap is shown in blue. The panels on the left and right columns show the sequence that corresponds to electron trapping via beam-capture and electron accumulation, respectively (see section C.1). A drawing of the cylindrical electrodes of the trap are added at the top of the columns. Two example images of electron clouds obtained via each of the two trapping techniques are included at the bottom.
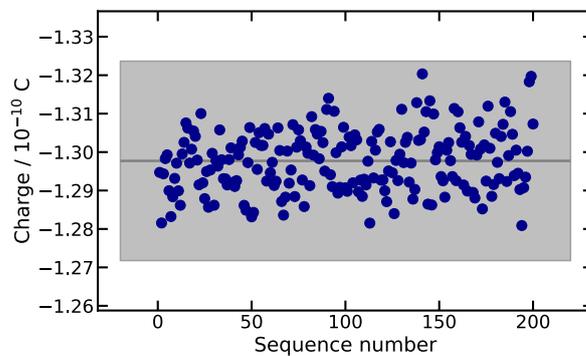
Figure C.2: The charged collected by the phosphor screen for 200 identical capture-store-release sequences with electrons trapped by the dynamic technique. The grey area shows the $\pm 2\%$ interval centred around the average collected charged.
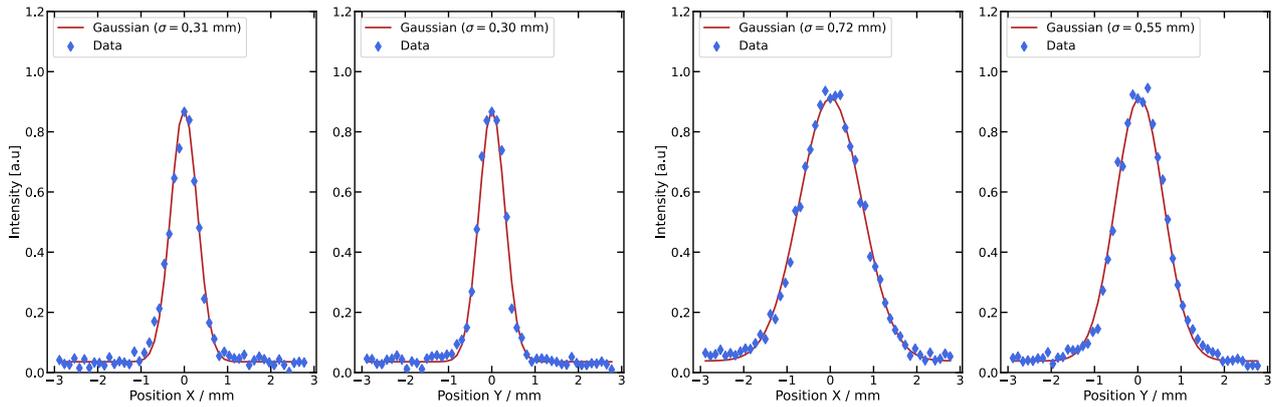
Figure C.3: The transverse size of two confined electron columns along the X and Y axes inside the trap as measured from releasing the electrons towards the phosphor screen. The electrons were trapped via the dynamic technique from the beam produced by the source with (left) and without (right) an additional collimator (the Wehnelt cylinder).

supplied to the end electrodes. As a consequence, the beam was typically captured in about $40\,\mu$s.

This rapid beam-capture technique is valuable for studies of the magnetron rotation of the electrons inside the trap. Firstly, the transverse size of the trapped electron cloud is comparable to the diameter of the beam, so that thin electron columns can be studied. Secondly, the initial transverse position of the electron cloud is determined by the radial off-set of the beam and can be changed by translating the position of the source. Hence, electrons can be captured at various radial locations with respect to the symmetry axis of the trap. Finally, due to the stability of the beam current and the fast rising of the trap potential, the dynamic technique leads to percent level variations in the total trapped charge between separate capture-store-release sequences, as shown in fig. C.2.

In the case of the beam-capturing technique, the total number of electrons that can be trapped is limited by the maximum beam current produced by the source. Similarly, the radial size of the electron column that becomes confined is dictated by the transverse size of the electron beam. Figure C.3 shows some typical transverse sizes of the electron column inside the particle trap. Wider electron clouds were produced by the source in the absence of the Wehnelt cylinder that acted as an additional narrow collimator. Oval shaped electron columns were frequently observed. An overview of transverse size and number of particles of the typical electron columns that were mainly confined to study the magnetron motion is shown in fig. 5.15.

Static particle trapping

An alternative technique to the rapid capture of the beam is based on the slower accumulation of electrons over a longer time interval, typically of tens or hundreds of milliseconds. The main steps in the accumulation sequence are schematically shown in fig. C.1. Before the accumulation, the electron beam propagates from the source to the grid in front of the MCP/P-screen detector. The first step of the particle trapping is to switch on the confining potential rapidly. However, in comparison to the dynamic particle trapping, the trap potential is
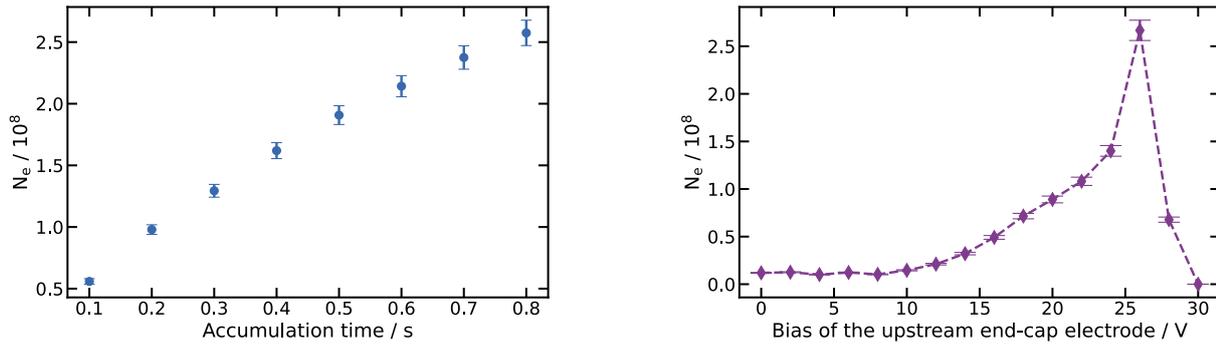
Figure C.4: The total number of electrons accumulated via the static trapping technique (electron accumulation). **Left:** The number of trapped electrons increases for longer time intervals during which electrons are accumulated. **Right:** For a fixed accumulation time interval, there is an optimal bias applied to the upstream end-cap electrode which maximises the total number of electrons that become trapped. The beam energy was $30\,\text{eV}$ and the outermost electrodes were used as end-caps.

asymmetric regarding the bias voltages applied to the two end caps. The upstream end electrode closer to the source is maintained at a potential which is a few Volts lower than the equivalent maximum potential of the beam (maximum parallel energy divided by the electron charge, e). The downstream end electrode of the trap is maintained at a high potential that fully blocks the electrons, typically a few tens of Volts above the maximum beam potential. The most energetic electrons from the source with an energy above the upstream blocking potential enter into the trapping volume and are reflected back by the downstream high-potential barrier. The electrons that enter the trap loose part of their energy through multiple collisions with the residual gas or through the two-stream instability. If enough energy is lost, the electrons cannot pass the entry potential barrier anymore and become confined.

The axial potential well can be kept 'open' for a long time interval of up to seconds in order to accumulate more electrons. However, as the space-charge density inside the trap increases and electrons that have already been captured spend longer times inside the trap, the rate of accumulation will eventually be compensated by the rate of loss of electrons due to radial outward transport. Below this threshold, the total number of accumulated electrons was observed to increase linearly with the time for which electrons were allowed to enter the confining volume, as exemplified in fig. C.4. Long accumulation times also lead to significant radial expansion of the trapped electron cloud. The height of the entry barrier of the potential well with respect to the maximum electron energy determines the energy spread of the particles that become confined. In this work, it was found that for an electron-beam energy of $30\,\text{eV}$, the largest number of electrons is accumulated for a fixed accumulation time when the bias of the electrode at the entry of the trap was close to $25\,\text{V}$ (see fig C.4). Raising the entry potential barrier closer to the maximum beam potential decreases the accumulation rate, but results in a lower energy spread of the confined particles. At the end of the accumulation interval, the voltage on the upstream end electrode is raised to block any further electrons from entering the trap. The space-charge is confined in a symmetric potential well for the desired storage time. Figure C.5 shows an example of the lifetime of an electron
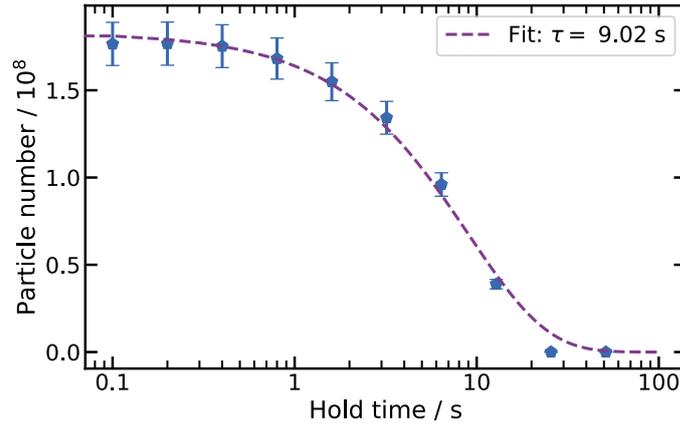
Figure C.5: The total number of electrons that remain confined as a function of the time interval during which they are hold within the trap. An exponentially decaying function is fitted to the data to extract the lifetime τ.

cloud trapped via the static accumulation which was typically several seconds.

C.2 System control

Some of the capabilities of the control system that were used in this work are: configuration of the data acquisition parameters and the triggers of the CCD camera and the digitiser, definition of the experimental sequences for capture, storage/study and ejection of the electron clouds, creation of experimental variables that changed during automated repeats of the same sequence, real-time visualisation of the data recorded by the MCP/P-screen assembly. The following outlines some of the components of the control system that allowed the efficient configuration of the experimental sequences, the actual run of the sequences, and the data logging.

Sequencer

The sequencer is the main virtual instrument that allows the user to configure an experimental run. A file is generated by the sequencer that initialises any variables that change during multiple repeats of the same measurement, generates analogue and digital outputs, and triggers the data acquisition. Running the file generated by the sequencer drives the experimental apparatus itself.

One of the main roles of the sequencer is to control the voltages that are supplied to the electrodes of the electron trap. This is done through an FPGA and two Analogue Output devices. There are 16 analogue outputs (AO) available which can each produce an output in the range $\pm 10\,\text{V}$. The AOs are connected to an amplifier with a gain of 14 to supply the trap electrodes with voltages in the range $\pm 140\,\text{V}$. Changes to these voltages require a minimum step duration of $3\,\mu\text{s}$. Faster digital outputs are also available with a minimum time step of $37.5\,\text{ns}$. Both the digital and the analogue outputs can be adjusted in steps of $12.5\,\text{ns}$ above their respective minimum values. Furthermore, the sequencer allows the definition of a number of steerable experimental parameters (termed VARs) which may be configured prior to running the FPGA.

A second capability of the control system is to produce a batch of individual sequences. An additional virtual instrument, the 'Modder', takes a template sequence file and replaces a number of features (e.g. time step duration, AO, VARs) with all possible permutations of the values given as an array list or generated from a formula provided by the user. The Modder creates the batch of individual sequences and a list instruction file that tells the Sequencer to execute each sequence in turn back-to-back. This capability has been extensively used in the present work, for example, to image the radial location of the electron cloud inside the trap while scanning the confinement time in equal steps over the range of interest. Thus, the range of confinement time intervals was provided as a set of equally spaced values and a batch of sequences were generated and run automatically, one for each duration of the confinement stage. Processing the recorded data from a batch of sequences provided temporal information about the evolution of the electron cloud within the storage trap.

C.3 Data recording

The acquisition devices that have been predominantly employed in the experiments with confined electrons described in this work are outlined below.

Digitiser

The digitiser was used predominantly to record the voltage response of the charge collector. The voltage signal was subsequently converted into charge to obtain the total number of electrons in the trap. The apparatus can record voltage readings at a maximum sample rate of two gigasamples per second across two separate channels, with 8-bit resolution in a $100\,\mathrm{mV}$ to $10\,\mathrm{V}$ range. The recording of the data by the digitiser is synchronised with the experimental sequence by a hardware trigger implemented on an FPGA.

The digitiser card is configured using Labview™ which extracts and saves the data on a disk storage. The resulting files are structured to contain the timestamp of the run and any sequences within the run and the settings of the digitiser, such as sampling rate and voltage range. The recorded files can subsequently be associated to a specific experimental run by their unique time identifiers. A set of Python scripts was used to extract the raw data from the structured files, display the voltage-time traces, and process parts of the signal. The scripts also allowed the fitting of functions to the data and automatic calculation of averaged quantities from multiple repeats of the same experimental run. Figure C.7 shows an example of a digitiser trace from the MCP/P-screen detector that resulted from the collection of an electron cloud released from the storage trap.

CCD Camera

The scintillation light from the MCP/P-screen assembly is imaged using a *Hamamatsu* Orca-R2, 1.3 megapixel CCD camera. The camera is configured using Labview™ which also extracts and saves the image data into a file structure consistent with that produced by the data acquisition from the digitiser. A hardware trigger on an FPGA ensures the synchronisation between the camera acquisition and the sequencer. The stored images

can then be paired with the corresponding voltage trace from the digitiser by the unique time and sequence identifiers. Various Python scripts were used to visualise the recorded images, overlap multiple images to show the time evolution of the electron clouds, subtract the nonuniform noise background across the CCD, and to fit a 2-D Gaussian function to the spatial distribution of the collected electrons. An example of the recorded image of an electron cloud is given in fig. C.11.

As outlined above, both the digitiser and the camera can be triggered during a sequence at precise points that are specified by the user from the sequence-control software through individual digital outputs. In most of this work, the camera and the digitiser were triggered when the axial potential barrier of the storage trap was lowered and the electrons were, thus, ejected from the confinement section of the beam-line towards the imaging detector.

C.4   Data processing

One key quantity that was derived from the measurements in this work is the total number of electrons confined within the plasma trap. The total number of electrons directly establishes the confinement capabilities of the trap and determines whether the dynamics of the trapped-electron ensembles that were studied correspond to the plasma or the single-particle regime. In order to derive the total number of electrons from the raw data, it was required to firstly measure the capacitance of the electron collector assembly and the amplification factor (gain) of the MCP.

Electron collector capacitance

The amplified charge collected by the phosphor screen or the unamplified charge picked-up by the front plate of the MCP is calculated via the relationship

$$V = \frac{Q}{C},$$
(C.1)

where V is the voltage signal measured by the digitiser when the electrons hit the MCP/P-screen assembly, Q is the calculated charge, and C is the capacitance of the collector. Thus, it was necessary to measure the effective capacitance in the two cases. Firstly, the capacitance of the MCP/P-screen assembly was needed when the P-screen was used as a collector and the MCP bias voltage was on. Secondly, a different balue for the capacitance was also required when the front plate of the MCP acted as the collector with no amplification provided by the MCP.

In both cases, the capacitance of the collector was calculated from a measurement of the time constant of the resonant circuit of which the collector is part. When the collector is hit by electrons, it acts as a discharging RC circuit with a known resistance $R = 1\,\text{M}\Omega$. The exponential decay of the charge in the collector circuit is equivalent to the exponential decay of the voltage recorded by the digitiser. An example of the time dependence of the voltage recorded by the digitiser before and after electrons hit the collector is shown in fig. C.6. As the
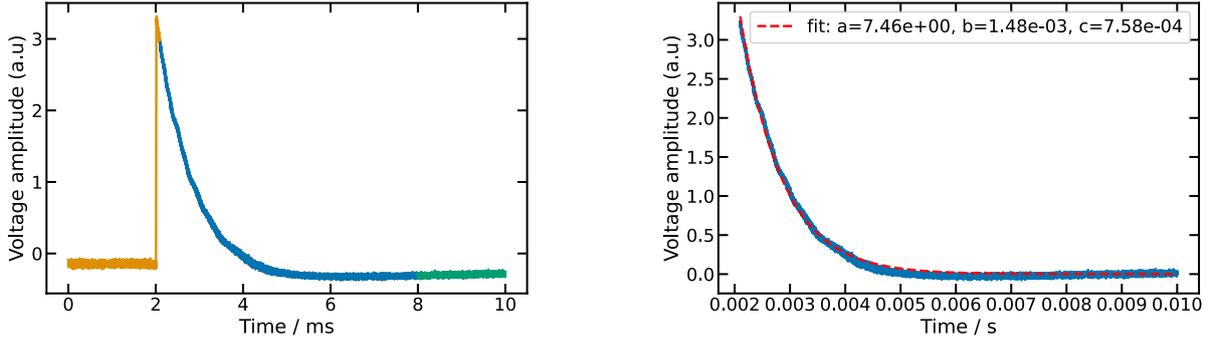
Figure C.6: **Left:** The voltage recorded by the digitiser when electrons are collected by the phosphor screen with the MCP turned on. The electrons hit the MCP/P-screen assembly at around 2 ms. **Right:** The decaying part of the signal shown in blue on the left is fitted with an exponentially decaying function to extract the time constant of the resonant circuit, and, hence, the capacitance $C = c \times 10^6$ pF.

Table C.1: Measured values for the capacitance of the MCP/P-screen assembly with and without amplification.

Collector	Amplification	Capacitance
MPC front plate	no (MCP bias off)	$697 \pm 7\%$ pF
P-screen	yes (MCP bias on)	$759 \pm 4\%$ pF

discharge of the resonant circuit corresponds to a voltage decay given by

$$V = V_0 e^{-t/RC}, \tag{C.2}$$

the decaying part of the signal recorder by the digitiser was fitted with the function $a \exp\left[-(t - t_0)/(RC)\right]$ to extract the capacitance C. The signal from the digitiser when electrons were collected was recorded several times in different days and for various total number of electrons collected. An average capacitance was then calculated from all the available measurements. The results are presented in table C.1 for the two cases of using the MCP front plate (no amplification) or the P-screen (MCP bias on, with amplification) as the electron collector.

It must be noted that the collector was used to measure the number of electrons that were ejected from the trap and travelled downstream to the MCP/P-screen assembly. It is assumed that the number of electrons lost while releasing them from the trap or within the beam-pipe between the trap and the imaging assembly is negligible. The efficient transport of the electrons is ensured by the coils positioned along the beam-line which generate a focusing magnetic field which, in turn, keeps the electrons away from the inner walls of the vacuum chambers.
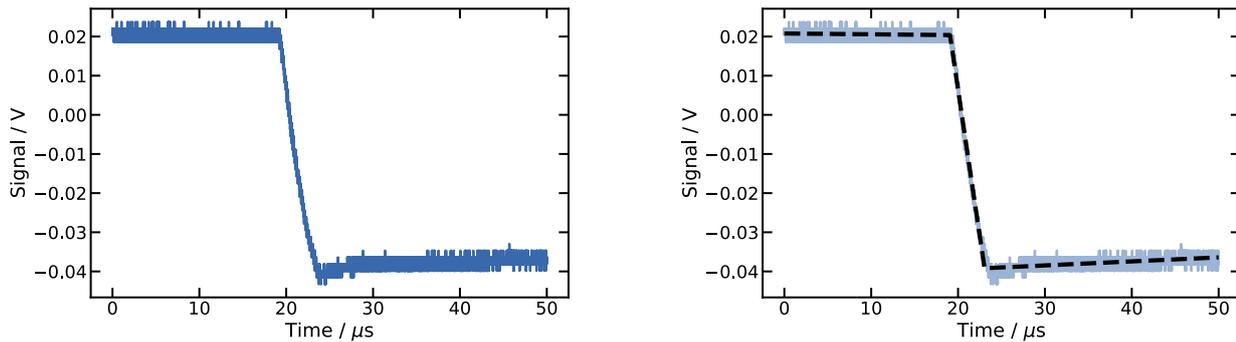
Figure C.7: An example of a voltage trace recorded by the digitiser when an electron cloud hit the MCP/P-screen collector. The signal data is recorded starting about 20 μs before the charge is collected. The raw signal on the left is fitted with a three-segment piece-wise linear function as shown on the right.

Analysis of the digitiser traces

When electrons are collected by the MCP/P-screen assembly, the digitiser records the voltage-time trace generated by the electrons rapidly charging up the collector. The voltage trace is processed to extract the rapid voltage change when the electrons hit the detector. The change in voltage is then converted into the total charge collected or the corresponding total number of electrons.

A typical voltage trace recorded by the digitiser is shown in fig. C.7. Typically, the digitiser was configured to start recording the voltage before the electron signal with at least 20% of the total acquisition time. The flat trace before the electron signal is used to calculate the average level of noise. When the charged particles hit the collector, the voltage drops rapidly with a typical overshoot due to the electronics. After the overshoot, the collector discharges on a much longer timescale compared to the time interval over which the trace is recorded.

In order to extract the change in voltage when the electrons hit the collector, the raw trace is fitted with a three-segment piece-wise linear function. An example of a fitted signal is shown in fig. C.7. The trace is analysed with a Python script and library [302] that finds a linear function with two breaking points that minimises the total sum of squared differences between the data points and the fitted function. Subsequently, the change in voltage generated by the collected charge is calculated as the difference between the mean value of the recorded signal before the first break point and the mean value of the signal after the second break point. In this case, the mean value of the signal after the second break point was used in the calculation to correct for the overshot of the electronics with the use of the positive slope of the signal after the second break point.

MCP gain calibration

The images recorded by the CCD camera were always obtained with the multi-channel plate turned on (both the front and back plates biased to certain voltages). Thus, the electrons that were ejected from the storage trap and hit the front plate of the MCP were amplified in order for a high enough charge to hit the phosphor screen and produce measurable scintillation light. Since the MCP bias was turned on when the phosphor screen was
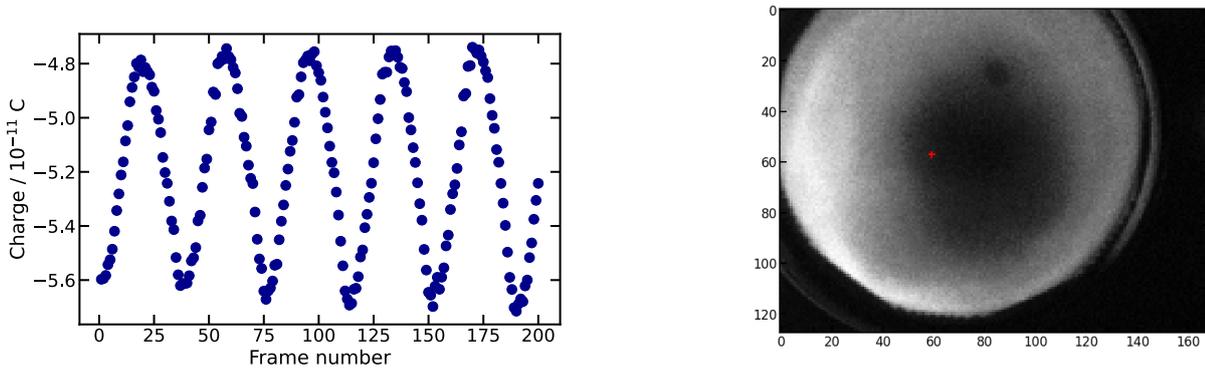
Figure C.8: **Left:** The variation of the charge amplified by the MCP and collected by the phosphor screen for multiple capture-store-eject sequences with an increasing storage time interval. Between different frames, the spot where the electrons hit the MCP rotates across the surface of its front plate, suggesting a non-uniform aplification factor. **Right:** Camera image of the light from the phosphor screen after a large electron cloud hit the MCP. The image shows the damaged areas on the surface of the phosphor screen.

used as the charge collector, the total charge recorded by the digitiser corresponds to the amplified number of electrons and not to the number of particles released from the plasma trap. As a result, the amplification factor or gain of the MCP must be known in order to convert the amplified charge recorded by the digitiser into the number of electrons that originated from the storage trap and were collected by the front plate of the MCP.

The MCP/P-screen assembly that was available at the time of this work had been previously used for experiments with positrons and electrons. It was observed throughout the measurements described in this work that both the amplification factor of the MCP as well as the amount of light emitted by the phosphor screen showed significant variations depending on the region where electrons hit the detector. Hence, the gain of the MCP was measured to be spatially nonuniform as a possible result of some of the channels of the MCP being damaged. Figure C.8 shows an example of the variation of the amplified charge measured as a function of time when a similar number of electrons were ejected from the storage trap at different radial positions on the front plate of the MCP. Furthermore, fig. C.8 shows areas of the screen which were damaged or less sensitive. These areas could be identified by dumping a radially wide electron cloud onto the MCP such that light was emitted by most of the active surface of the phosphor screen.

A calibration of the gain of the MCP was carried out to determine the average amplification factor and to evaluate its spatial variation. The uncertainty in the amplification factor of the MCP dominates the errors in the calculations of the total number of confined particles throughout the measurements described in this chapter. The amplification of the MCP was calculated from repetitions of the same two types of measurements performed with identical electron capture-storage-release sequences. Firstly, the collected charge was measured with the amplification of the MCP turned on and with the P-screen as a collector. Secondly, the same sequence with an approximately identical number of trapped electrons, was run to measure the charge collected by the front plate of the MCP with the MCP amplification turned off. In the case of the second type of measurement, the signal from the charge collected by the MCP was amplified by an external amplifier with a known gain. As a result of each pair of measurements, the gain of the MCP was calculated as the ratio between the charge collected
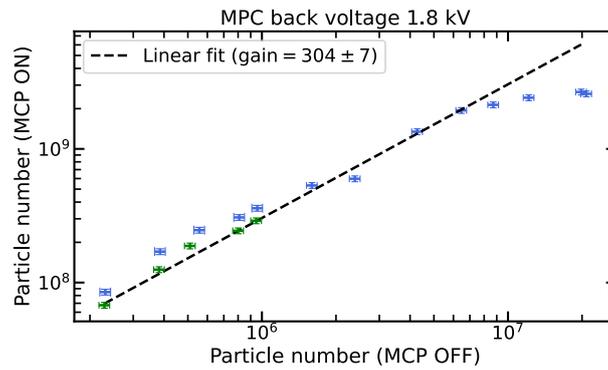
Figure C.9: The total electron number measured with the MCP bias voltages on versus the total electron number measured with no amplification from the MCP. The gain starts to saturate above about 10^7 incident particles. A linear fit is applied to the data points below saturation to extract an average amplification factor. The data points have two distinct colours to indicate that they were taken five months apart.

following the amplification of the MCP and the non-amplified charge.

The pair of measurements described in the previous paragraph was repeated for several electron-beam currents from the source. Figure C.9 shows how the charge amplified by the MCP changes as a function of the total charge incident on the front plate of the MCP over a wide range of total number of electrons released from the storage trap. A linear fit was applied to extract the average amplification factor. A saturation of the amplification capability of the MCP was observed for large numbers of electrons hitting the detector. The data points that indicate no further increase in the gain were excluded from the linear fit. Moreover, fig. C.10 shows several measurements of the amplified charge as a result of the same capture-storage-release sequence for a few distinct release times. As a delay in the release time leads to the electron column rotating further inside the storage trap, the electrons ended up hitting the MCP at distinct locations. These measurements were used to calculate the maximum and minimum amplification factor across the surface of the MCP over which the electrons were collected. It must be noted that this procedure did not scan the full surface of the MCP front plate so it does not ensure the existence of regions with an amplification factor outside of the range measured with the current procedure. Nonetheless, the electron source was positioned such that the charge hit the MCP close to the regions of maximum and minimum amplification as observed on the images of the P-screen. Hence, the variation of the measured amplification factor should be close to the actual gain provided by different regions of the MCP.

Throughout the measurements presented in this work, the minimum and the maximum amplification factors were measured to be 207 and 548, respectively. Given the large variation of the amplification factor of the MCP with the region where the electrons hit the detector and the scan of gain factors shown in fig. C.9, the gain used in the results presented below is considered to be $300 \pm 50\%$.
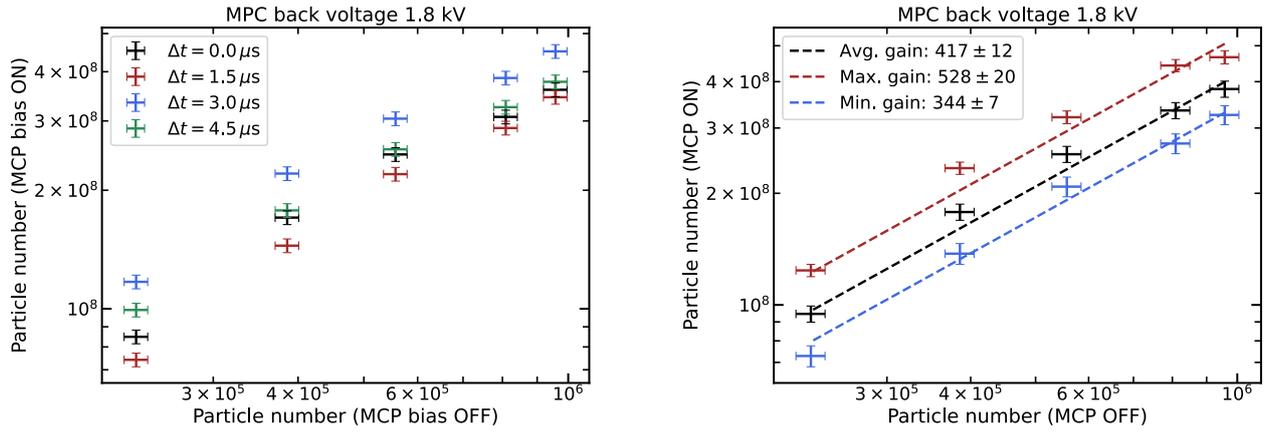
Figure C.10: The total number of electrons collected with the amplification of the MCP turned on as a function of the total number of electrons measured with no amplification from the MCP. **Left:** The electrons were released from the storage trap with four distinct time delays such that they hit the MCP at different locations. **Right:** The minimum, average, and maximum particle number that were measured during the rotation of the spot where the electrons hit the MCP. A linear was applied to extract the corresponding amplification factor of the MCP.

Image processing

Once the capacitance and gain of the MCP/P-screen detector were determined, the rest of the data analysis consisted of image processing. The images recorded by the CCD camera were passed through several steps of processing to extract information on the transverse distribution of the confined electron clouds typically as a function of the time duration for which the electrons were trapped with respect to a reference time. Figure C.11 summarises the most important image-processing steps which are detailed below. As the majority of the results of this work involve the magnetron rotation of the electrons inside the storage trap, the steps outlined below describe how the period and amplitude of magnetron motion were computed from the data.

Each experimental run typically consisted of one hundred individual capture-store-release sequences. Each sequence differed from the previous one by a fixed time delay for which the electrons were trapped during the storage stage of the sequence. The time delay was measured with respect to the storage time used in the first sequence from the full batch of sequences. Releasing the electron column towards the phosphor screen at equally spaced points in time is equivalent to recording the transverse position of the electrons at distinct equally spaced phases during their magnetron rotation. Each image of the phosphor screen recorded by the camera was then processed separately.

Firstly, the background noise was removed from each image. Twenty images were taken in the absence of any electron signal and the average background noise was calculated. The average noise was subsequently subtracted from the images that contained an electron signal. Then, a 2-D Gaussian fit was applied to the intensity profile to extract the position of the peak, (x_0, y_0), the rotation angle of the symmetry axes of the distribution with respect to the X–Y coordinate system, θ, and the radial size of the electron cloud on the screen, σ_x, σ_y.
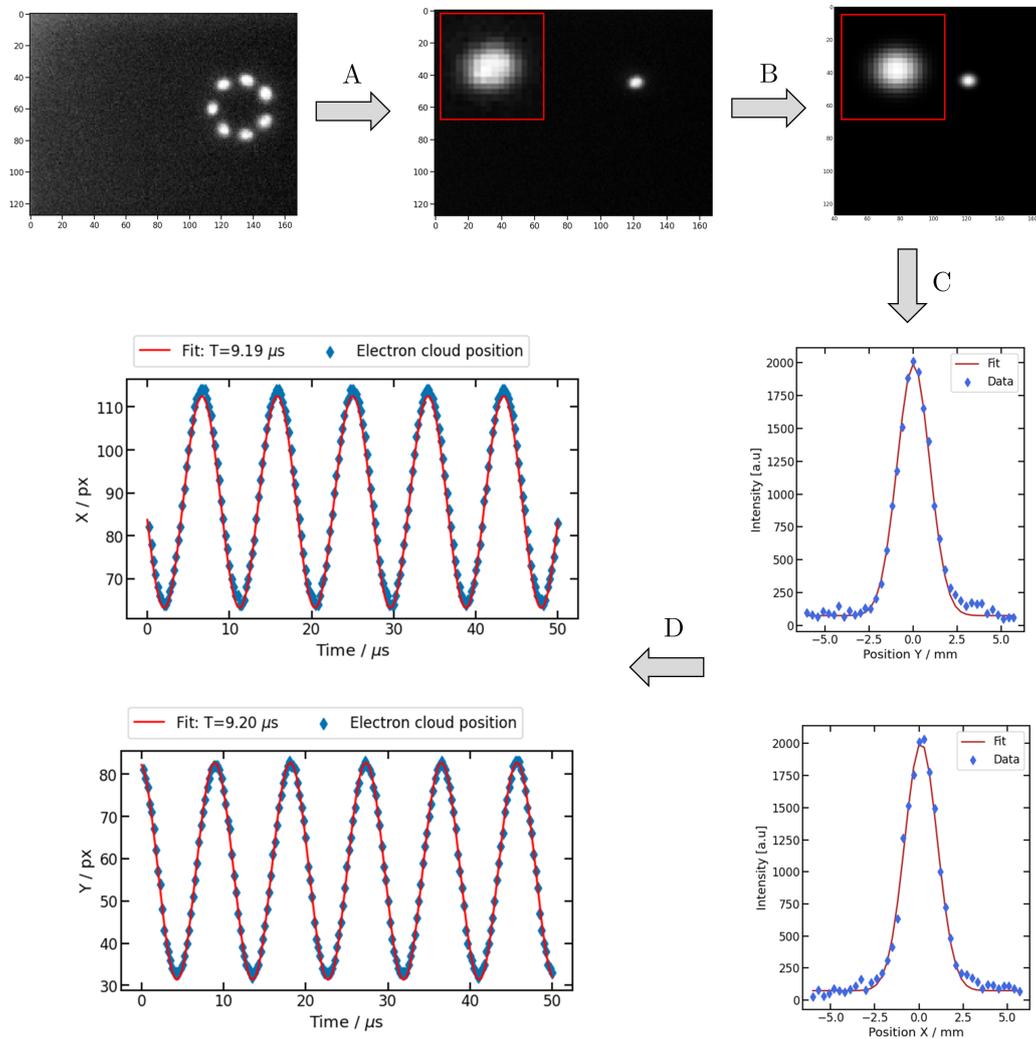
Figure C.11: Summary of the typical steps employed in processing the images from the phosphor screen. **A**: An experimental run is recorded as individual camera images for each separate capture-store-release sequence in the run. **B**: The background noise is removed from each recorded image and a 2-D Gaussian fit is applied to the remaining signal. **C**: The radial size of the electron column on the phosphor screen is computed from the projection of the 2-D Gaussian fit onto the X and Y axes. **D**: The resulting position of the electron cloud is computed for each individual experimental sequence and plotted as a function of the time duration during which electrons are confined. A sinusoidal function is separately fitted to both the X and Y data points to extract the period and amplitude of the electron cloud rotation. The two insets on the top row show a magnification of the high-intensity spot in the corresponding image.

The Gaussian function used for the fitting procedure is

$$f(x,y) = A + B \exp\left[a(x-x_0)^2 + b(x-x_0)(y-y_0) + c(y-y_0)^2\right], \tag{C.3}$$

where A, B are constants and $a = \cos^2\theta/2\sigma_x^2 + \sin^2\theta/2\sigma_y^2$, $b = -\sin^2(2\theta)/4\sigma_x^2 + \sin^2(2\theta)/4\sigma_y^2$, $c = \sin^2\theta/2\sigma_x^2 + \cos^2\theta/2\sigma_y^2$.

Subsequently, the position of the peak, (x_0, y_0), was plotted separately for the two axes as a function of time. A sinusoidal function was, then, fitted individually to the two data sets to compute the amplitude and period of the transverse rotation of the electron cloud. It was observed that most of the highly compressed clouds resulted in good Gaussian fits to the light-intensity distribution recorded by the camera. The Gaussian profile was well maintained during the magnetron motion for up to tens of complete rotations. For rotation at large radii, a low intensity tail similar to that of a comet was observed to appear in addition to the Gaussian peak.

The rotation amplitude and the radial size of the electron cloud on the screen were converted from pixels to millimetres using the factor $3.4\,\text{px/mm}$. The conversion factor was obtained by imaging the phosphor screen with the light tight box opened and measuring the aperture of the screen in both pixels and millimetres. The linear dimensions measured on the screen were further multiplied by a factor of 2.57 to obtain the actual radial offset and size of the electron cloud in the storage trap. This factor is applied to account for the magnetic-mirror effect created by the different magnitudes of the magnetic field within the trap and at the location of the phosphor screen. As the trajectories of the electrons are tightly bound around the magnetic field lines, the transverse size of the electron ensemble changes when the particles propagate through regions with different density of the magnetic field lines.

Appendix D

The diocotron instability

Stability of the diocotron waves

The stability of the diocotron mode is strongly dependent on the radial electron density profile. A uniform density distribution corresponds to rigid rotation of the plasma column which is always stable. Any radial variation of the plasma density causes a shear in adjacent azimuthal layers of the plasma which may lead to the formation of an instability (see fig. D.1). The shear is expected to drive the plasma to evolve into a rigid rotor steady state [232]. However, the evolution is much slower than the typical plasma oscillations and wave phenomena. For linear modes, it was shown that a necessary condition for the onset and growth of an instability is that $\partial n_e/\partial r$ changes sign over the interval $0 \leq r \leq b$ where b is the radius of the conducting cylinder within which the plasma is confined [190] (i.e. the radial plasma density profile is non-monotonic). Equivalently, a sufficient condition for stability for small-amplitude electrostatic flute perturbations (i.e. no longitudinal gradients, $\partial/\partial z = 0$, of plasma density or potential) is that the radial density profile, $n_e(r)$, is a monotonically decreasing function of radius r [231].

A radially-inverted electron population is described by the general simplified density profile shown in Fig. D.2 which corresponds to the step-function

$$n_e(r) = \begin{cases} \Delta \cdot n_e = \text{const.}, & 0 \leq r < r_b^-, \\ n_e = \text{const.}, & r_b^- < r < r_b^+, \\ 0, & r_b^+ < r \leq b. \end{cases} \tag{D.1}$$

Here, the parameter Δ varies from $\Delta = 1$ (uniform density distribution) to $\Delta = 0$ (annular density profile), the plasma column becoming strongly unstable towards lower values of Δ. The proximity of the outer conducting wall at radius b also influences the stability.

A density profile such as the one described by (D.1) offers a first-order approximation to some of the electron density distributions observed in numerical simulations. Furthermore, there is theoretical reason to expect a hollow electron cloud when the external magnetic field is larger than the value required to confine a specific electron density [47]. At stronger magnetic fields, the gyration radius of the electrons becomes smaller and the
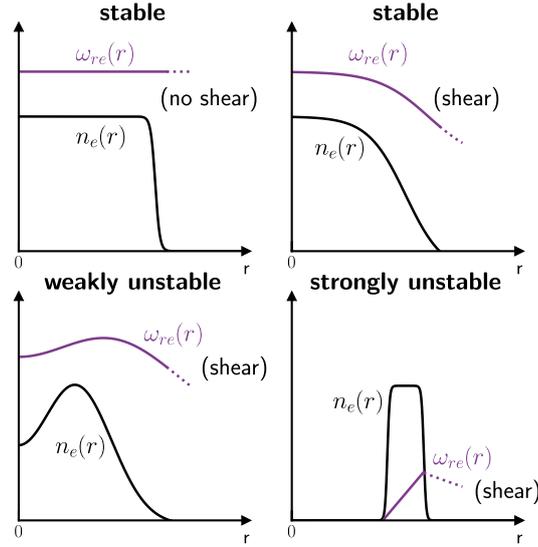
Figure D.1: Radial density distributions and corresponding angular density profiles for cylindrically symmetric long plasmas. The monotonically decreasing profiles shown on the top row corresponds to stable plasma states, while the hollow profiles shown on the bottom row can lead to azimuthal instabilities. Reproduced from [232].

electron column is compressed. Thus, the local electron density exceeds the confinement limit given by the anode potential and electrons escape from the central region.
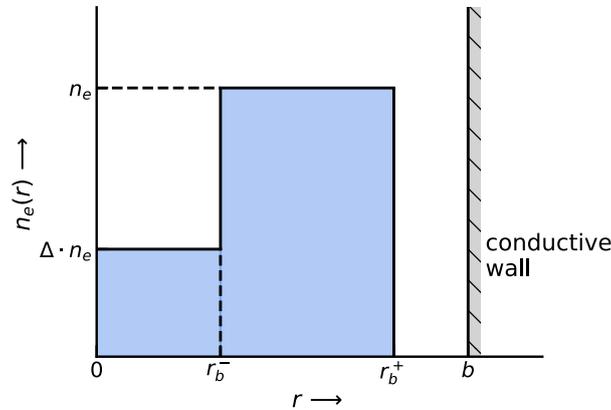
Figure D.2: Radial electron density $n_e(r)$ for a step-function profile with inverted population and a degree of hollowness Δ inside a conducting cylinder of radius b.

A linear stability analysis of the diocotron modes was derived in [231] for both the electron density profile drawn in Fig. D.2 and for a more general continuously-varying profile. Starting from a self-consistent Poisson equation, an eigenvalue equation is derived for low-frequency electrostatic flute perturbations (perturbations extended along magnetic field lines). For the step-function density profile in eq. (D.1) and the appropriate boundary conditions the following dispersion relation is obtained [231]

$$\left(\frac{\omega}{\omega_D}\right)^2 - 2\hat{b}\left(\frac{\omega}{\omega_D}\right)^2 + \hat{c} = 0 \tag{D.2}$$
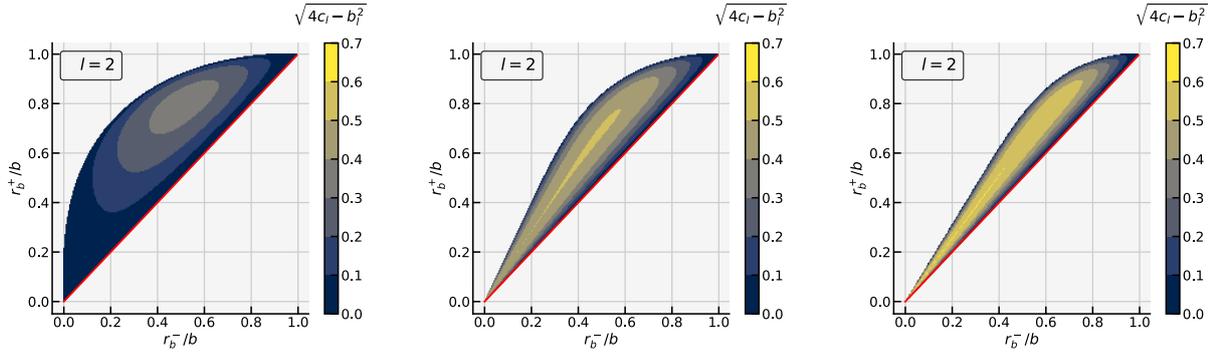
Figure D.3: Instability regions of the diocotron mode l for an annular electron density profile ($\Delta = 0$) extending from radius r_b^- to r_b^+ inside a cylindrical conductor of radius b. A higher geometrical coefficient $\sqrt{4c_l - b_l^2}$ corresponds to a higher instability growth rate ($b_l \equiv 2\hat{b}$, $c_l \equiv \hat{c}$ for $\Delta = 0$). An unstable solution exists only in the region above the red line where $r_b^+ > r_b^-$. The system is always stable for the fundamental mode $l = 1$.

where the geometrical factors $\hat{b}$ and $\hat{c}$ are defined by

$$2\hat{b} = l[1 + \Delta - (1 - \Delta)(r_b^-/r_b^+)^2] - \{[1 - (r_b^+/b)^{2l}] - (1 - \Delta)[1 - (r_b^-/b)^{2l}]\} \tag{D.3}$$

and

$$\hat{c} = l^2[1 - (1 - \Delta)(r_b^-/r_b^+)^2]\Delta - \{l[1 - (r_b^+/b)^{2l}]\Delta - l[1 - (r_b^-/b)^{2l}](1 - \Delta)[1 - (1 - \Delta)(r_b^-/r_b^+)^2]$$
$$+ (1 - \Delta)[1 - (r_b^-/r_b^+)^{2l}][1 - (r_b^+/b)^{2l}]\}, \tag{D.4}$$

where $\omega_D \equiv \omega_{pe}^2/(2\omega_{ce})$ represents the diocotron frequency and l the instability mode number. The solution to the dispersion relation is given by

$$\omega = \omega_D \left[\hat{b} \pm (\hat{b}^2 - \hat{c})^{1/2}\right] \tag{D.5}$$

from which the necessary and sufficient condition for instability is

$$\hat{c} > \hat{b}^2 \tag{D.6}$$

Whenever the instability condition is met, the growth rate is $\mathrm{Im}(\omega) = (\hat{c} - \hat{b}^2)^{1/2}\omega_D$.

In the special case $\Delta = 0$, Fig. D.3 shows the instability regions of the diocotron modes $l = 2, 3, 4$ for an annular electron density profile. For a fixed value of r_b^-/b, a sufficiently large value of r_b^+/b leads to stability. Thus, the proximity of the outer conducting wall has a stabilising influence on the diocotron modes [190].

A different behaviour is seen for the more general case $\Delta \neq 0$. Fig. D.4 shows the stability-instability contours for the first three unstable diocotron modes l. The instability condition (D.6) is met inside the closed contours. As the electron density profile becomes less hollow, the parameter space corresponding to instability is reduced and shifted towards values of r_b^+/b closer to 1. As the mode number l increases, the instability regions move towards the $45°$ line.

Consequently, in practice, the lens has to be filled as homogeneously as possible in order to operate in a regime away from the diocotron instability. The homogeneity can be controlled to a certain degree by ensuring that the electrons are emitted from a cathode with a wide range of transverse canonical momenta. Separately, the axial magnetic field may lead to the formation of a hollow plasma column if not properly matched to the local electron density.

Negative energy of the m=1 diocotron mode

The simplest model of the $m_\theta = 1$ diocotron mode in a Penning-Malmberg trap is built on the assumption of an infinitely long plasma column which is displaced from the axis of symmetry of the trap. The motion of the electrons in the plasma are dictated by the external axial magnetic field and the electric field produced by non-neutral line of charge of the plasma. The interaction of the plasma column with its own electric field can be modelled by replacing the conducting wall of the trap with an equal and opposite image charge as in fig. D.5. The image charge is located at a radial distance S from the plasma column such that the electric potential is constant on the contour $r = R_w$, where R_w is the radius of the trap and r is measured from the axis of symmetry of the trap.

The requirement that $\phi(R_w, \theta) = \text{constant}$ is satisfied for a separation [303]

$$S = \frac{R_w^2}{D}. \tag{D.7}$$

The electrostatic energy of the mode is equal to the energy required to displace the plasma by a distance D in the electric field of the image charge,

$$W_{\text{ES}} = \int_0^D F dx = \int_0^D Q \cdot E_i dx = \int_0^D Q \cdot \left(\frac{-N_L q D}{2\pi\epsilon_0 R_w^2} \right) dx = -\frac{(N_L q)^2}{4\pi\epsilon_0} \frac{D^2}{R_w^2} L_p. \tag{D.8}$$

In simpler terms, the image charge has the opposite sign of the 'real' charge. Therefore, as the plasma is attracted towards its image charge, the electrostatic energy decreases as the diocotron mode amplitude increases. Here, the change in the kinetic energy of the plasma is ignored. The fact that the electrostatic energy of the plasma is negative implies that the mode can grow due to energy dissipation. This destabilisation of the diocotron mode is referred to as 'resistive growth' [304].
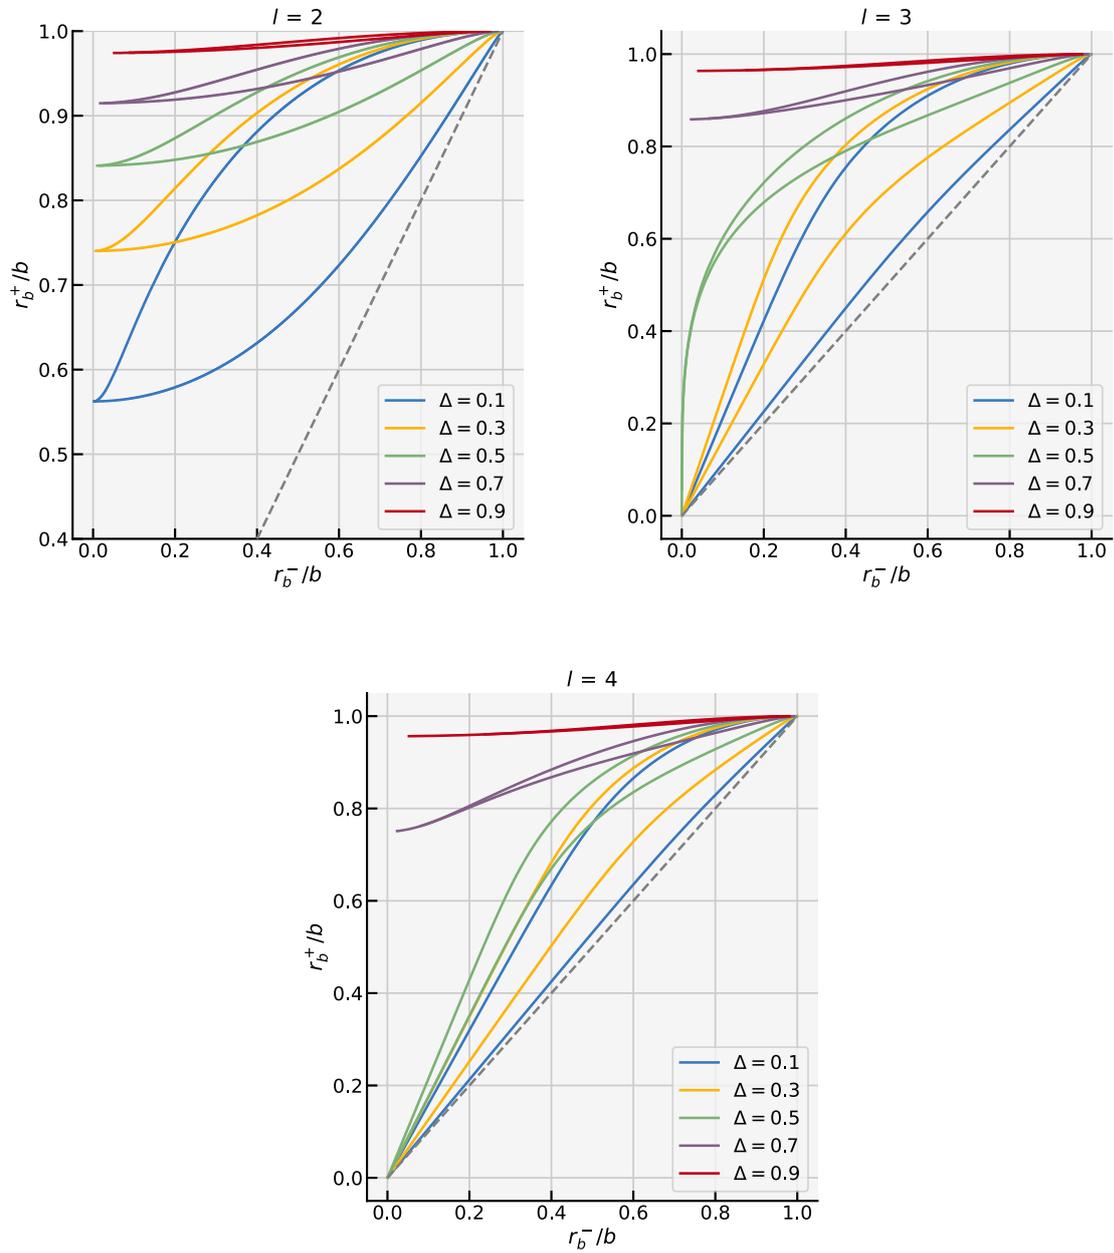
Figure D.4: Stability-instability contour plots of the diocotron mode l given by $\hat{b}^2 = \hat{c}$ as obtained from equations (D.3) and (D.4) for the step density profile described in Fig. D.2. Five filling factors $\Delta = 0.1, 0.3, 0.5, 0.7, 0.9$ are shown for each mode. The condition for instability is met in the regions inside the closed contours. A physical solution exists only above the dashed line where $r_b^+ > r_b^-$.
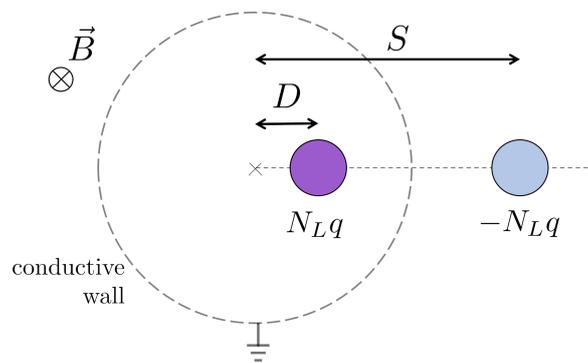
Figure D.5: Image charge model of the diocotron wave for an infinitely long plasma column in a Penning-Malmberg trap. Figure reproduced from [303].

Appendix E

Positron focusing by a confined electron plasma

This appendix describes a series of particle-tracking simulations that were carried out to assess the experimental conditions under which an electron plasma contained by one of the particle traps at Swansea University could produce a measurable focusing effect on a low-energy positron beam. The study was carried out prior to the preliminary measurements described in chapter 5. Given the existing layout of the beam-line, the purpose of the numerical studies was to determine the range of electron plasma densities that needs to be trapped in order to obtain a measurable focusing effect on a low-energy positron beam passing close to or through the plasma. Thus, this subsection presents a set of single-particle and Monte Carlo beam tracking simulations of positrons through a model of the electron plasma trap. The consequence of the axial magnetic field of the trap and the radially-focusing electric field of the plasma was found to be a rotation of the positron beam which is described in more detail towards the end of the subsection.

Due to the capability of the positron beam-line at Swansea to produce both positron and electron beams, as well as to trap both types of particles, a computer model of part of the beam-line was developed to investigate the focusing effect of an electron plasma on a low-energy positron beam and to assess the experimental conditions that could lead to a measurable focusing effect. It was found that **positrons with energies in the range $50\,\mathrm{eV}$–$100\,\mathrm{eV}$ that propagate through an electron plasma confined within a Penning-Malberg trap undergo an azimuthal rotation around the axis of symmetry of the trap**. The rotation is caused by the $E \times B$ configuration of the two fields: the axial magnetic field of the trap and the radial electrostatic attraction generated by the electron space-charge. If the space-charge of the electrons is uniform, the angular frequency of the positrons passing through the trap is independent of the radial position. The influence of the electron plasma on the positrons can be quantified by measuring the rotation of an off-axis pencil-beam or an on-axis elliptical beam of positrons around the axis of symmetry of the trap. **An azimuthal rotation of the positron beam that is measurable at the location of the detector downstream of the trap was found in particle-tracking simulations to be achieved for an electron plasma characterised by $4 \times 10^{12} \lesssim \mathrm{n_e L} \lesssim 4 \times 10^{13}\,\mathrm{m^{-2}}$.** Here, n_e is the density of the uniform electron plasma, L is the length of the plasma and the average axial magnetic field of the trap was approximately $40\,\mathrm{mT}$. However, it must be noted that in practice, a wide enough electron plasma must be trapped for the positron beam to pass through the quasi-uniform region of space-charge.

Furthermore, the electron plasma needs to be stored for a relatively long time that can allow the electron source to be retracted from the beam axis and for the positrons to be released towards the electron plasma.

Field maps for the positron beam-line

For the studies described in this section, the electron plasma was modelled as a cloud of uniform charge density that fills the internal volume of the trapping electrodes. In a real trap, the electron density is typically uniform in the bulk of the plasma and it falls to zero at the edges of the plasma over a distance of several Debye lengths. Thus, the assumption of uniform electron density leads to a valid model of the electric field in the region of the plasma close to the central axis of the trap. For a uniform space-charge, the radial component of the electric field, $\mathbf{E}$, is given by

$$E_r = -\frac{n_e e}{2\epsilon_0} r \tag{E.1}$$

where n_e is the electron density and r is the radial distance from the central axis of the trap. A 3D electric field map was generated from equation (E.1). Due to the uniformity of the space-charge along the z axis, and assuming a net zero effect of the longitudinal electric field component, E_z, the field map had $E_z = 0$ set throughout the entire volume. The electric field map was generated on a uniform 3D grid and a cubic interpolator was used by BDSIM. No difference could be observed when switching between the 'cubic' and the 'nearest' interpolators available in BDSIM. For particle-tracking inside the electron plasma, the maximum step size in the integration was restricted to $2\,\mu\text{m}$.

Special care was taken to define the field map in a volume that extends longitudinally past the entry and exit plane of the actual electron cloud. While BDSIM uses the field map for particle-tracking only inside the volume of the electron cloud, it was observed that sometimes the field map is not calculated on the entry and exit planes of the electron cloud if the longitudinal extent of the electric field map is the same as the length as the element that represents the space-charge cloud. To verify the tracking through the generated field map, the optical functions were compared with a separate simulation in which the space-charge cloud was modelled as a transfer matrix that corresponds to a thin lens with the focal length calculated from the electron density of the plasma. All the optical functions and the beam size showed good agreement between the two models of the plasma lens.

While an electron cloud with uniform density provides a radial focusing force for a positron beam, the presence of the axial magnetic field required to confine the plasma may modify the focusing effect due to the low energy of the positrons. In order to asses the effect of the magnetic field, the coils around the storage trap and the central solenoid have been modelled. The resulting magnetic field map was then used as input to the beam-tracking simulations.

A secondary focusing effect in a plasma lens is produced by the axial magnetic field, $\mathbf{B}$, that traps the electrons radially. Hence, it is necessary to account for the component of the Lorentz force generated by the

Table E.1: Specifications of the magnetic coils and solenoid modelled for the positron beam-line at Swansea.

Parameter	Solenoid	Steering coils
Internal diameter [mm]	146	526
External diameter [mm]	162	600
Length [mm]	680	60
I [A]	30	6
Turns	680	400
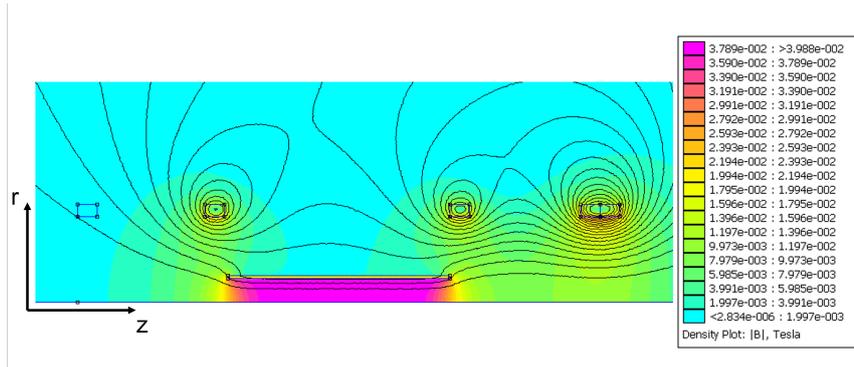
Figure E.1: The magnetic field strength and field lines in the (z, r) plane around the positron beam-line as obtained with FEMM [223]. The rectangular contours mark the position and size of the two pairs of quasi-Helmholtz coils and the main solenoid in-between.

magnetic field on the positrons for a realistic model of the lens. The magnetic field of the trap is produced by a water-cooled copper solenoid. Two pairs of steering (quasi-Helmholtz) coils are positioned around the two vacuum crosses immediately upstream and downstream of the solenoid (see fig. 5.2). The parameters of the main solenoid and the steering coils are given in table E.1. The coils and solenoid were simulated using a finite element analysis package FEMM [223]. The software is used to solve 3D axisymmetric linear magnetostatic problems with a low computational cost. For the studies presented here, FEMM was used to obtain a 2D, (r, z), axisymmetric magnetic field map starting from the value of the current through the coils and their configuration. The resulting field strength around the beam-line is shown in fig. E.1. The field map obtained with FEMM was then loaded into the particle tracking code BDSIM [141] with the help of a regular grid interpolator and the interface between Python and BDSIM. The trap electrodes are placed in the middle of the solenoid, in a uniform field region of about $400\,\text{G}$. The magnetic field outside of the electrodes is not zero due to the steering coils and the two end sections of the solenoid.

Single-particle trajectories

The first step in studying the focusing effect of an electron plasma on a low energy positron beam under the conditions of a real experiment was to track individual positrons through the electric field of the space-charge combined with the magnetic field generated by the coils around the beam-line. Several cases were studied for various electron densities and steering coils turned on or off. For each case, three positrons were simultaneously

tracked through the beam-line with initial offsets from the beam axis of 1, 2, and 3 mm, respectively. A kinetic energy of 85 eV was chosen for the positrons. The transverse divergence of the positron beam was fixed to an arbitrarily small value of $x' = y' = 10^{-6}$ rad.

Figure E.2 shows the axial magnetic field strength along the beam-line and the volume occupied by the electron plasma, together with the trajectories of three positrons with different initial transverse offsets. In the regions with only a magnetic field, the positrons undergo cyclotron rotation around the magnetic field lines. The gyroradius can be seen to decrease in the sections with a lower magnetic field strength. In the region occupied by the electron plasma, the $\mathbf{E} \times \mathbf{B}$ field causes an additional rotation of the guiding centre around the central beam axis, as shown schematically in fig. E.3.

Assuming that the gyroradius is small such that the changes in the radial component of the electric field can be neglected, the guiding centre of the cyclotron rotation will drift with a drift velocity

$$\mathbf{v_d} = \frac{\mathbf{E} \times \mathbf{B}}{B^2} \tag{E.2}$$

For a uniform electron cloud, $E \propto r$ implies that $v_d \propto r$. Therefore, the angular velocity of the positrons is independent on their radial position which results in the entire beam rotating around the beam axis as a rigid rotor. The angular velocity increases with the density of the plasma, as can be seen in fig. E.2 from the increase in the number of full rotations of a positron while passing through the electron cloud at a larger plasma density.

To verify the results from the single particle-tracking simulation, a separate script was written to integrate numerically the full equations of motion

$$\begin{cases} \ddot{x} = \omega_c \dot{y} - \alpha x \\ \ddot{y} = -\omega_c \dot{x} - \alpha y \\ \ddot{z} = 0 \end{cases} \tag{E.3}$$

where $\omega_c = eB/m_e$ and $\alpha = n_e e^2/2\epsilon_0 m_e$. Figure E.4 shows the trajectory obtained in the transverse plane for a positron passing through the storage trap. For the case of $n_e = 2 \times 10^{15} \, \text{m}^{-3}$ and $L = 19 \, \text{cm}$, integrating the radial acceleration numerically gives an average radial velocity of approximately $-0.7 \, \text{ms}^{-1}$. For a transit time of about $3.5 \times 10^{-8} \, \text{s}$, the numerical result is consistent with the absence of a significant radial drift of the guiding centre. In other words, the positrons are not focused towards the central axis when passing through the lens. Instead of an overall focusing effect, the positrons experience on average a rotation around the symmetry axis of the trap.

Evolution of the positron-beam size

In addition to single particles, a beam of positrons was also tracked through the beam-line. The initial beam distribution was defined to be Gaussian with $\sigma_x = \sigma_y = 3 \, \text{mm}$, $\sigma_{x'} = \sigma_{y'} = 10^{-6}$ rad, a kinetic energy of 85 eV, and no energy spread.

Figure E.5 shows how the size of the positron beam changes throughout the electron plasma and downstream
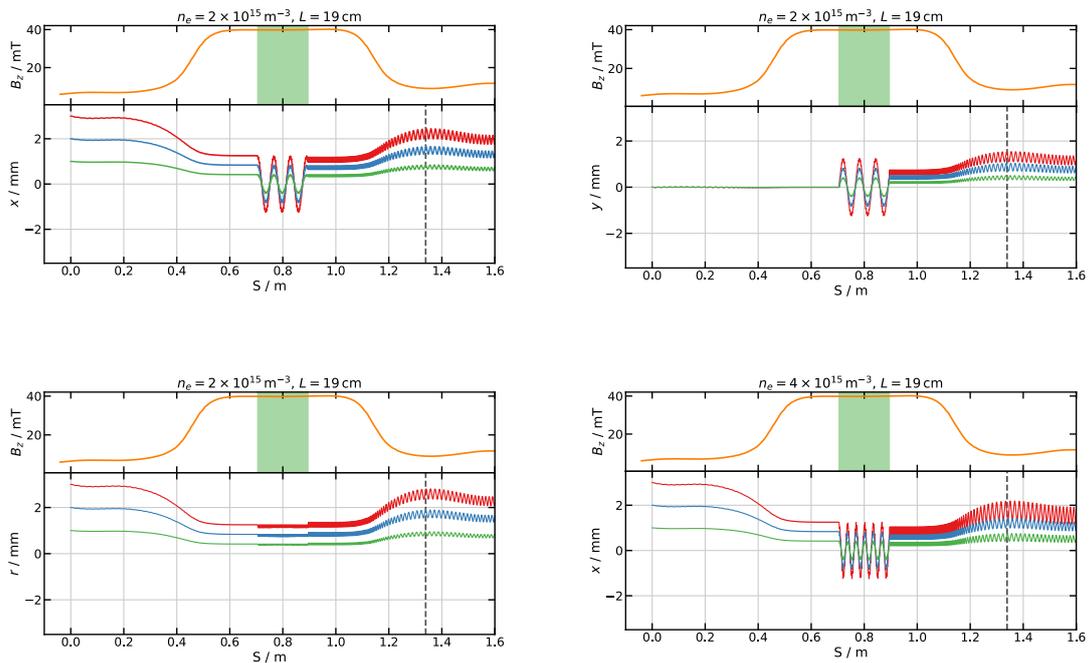
Figure E.2: Single particle trajectories for three positrons with initial offset of 1,2, and 3 mm from the beam axis. The top panel of each figure shows the axial magnetic field strength B_z along the beam-line and the region filled by the electron plasma in green. The motion in the transverse x and y directions, as well as the change in radius r show a cyclotron rotation of each positron around a guiding centre. The guiding centre itself rotates around the beam axis while passing through the electron plasma. The dashed line marks the position of the detector (MCP).
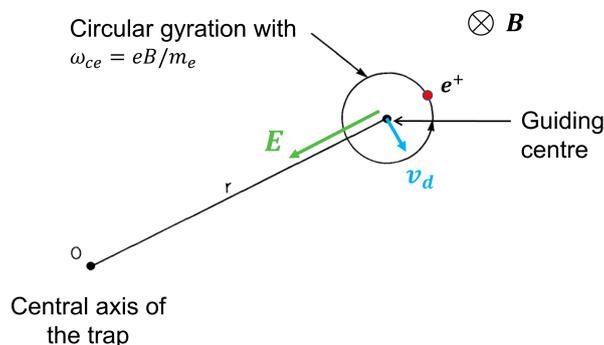
Figure E.3: Schematics of the trajectory of a positron in a $\mathbf{E} \times \mathbf{B}$ field. The particle undergoes cyclotron motion around the $\mathbf{B}$ field lines. In a constant radial electric field $\mathbf{E}$, the guiding centre of the circular motion drifts around the central axis with a velocity v_d.
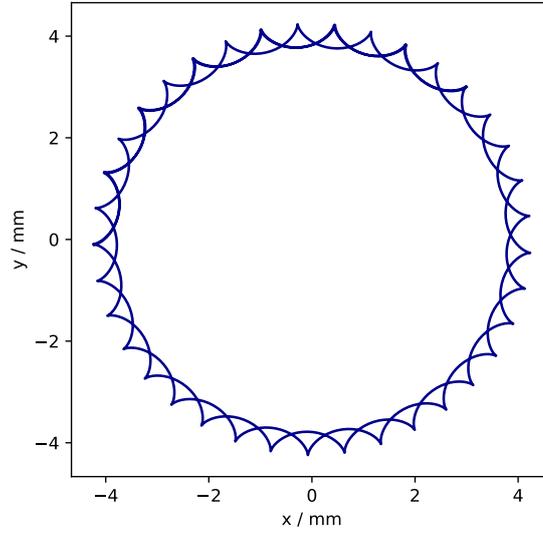
Figure E.4: Trajectory in the transverse plane for a positron that passes through the $\mathbf{E} \times \mathbf{B}$ field within storage trap as numerically calculated from equations (E.3). Due to the short transit time and small radial variations in $\mathbf{E}$, no radial drift of the trajectory can be observed.
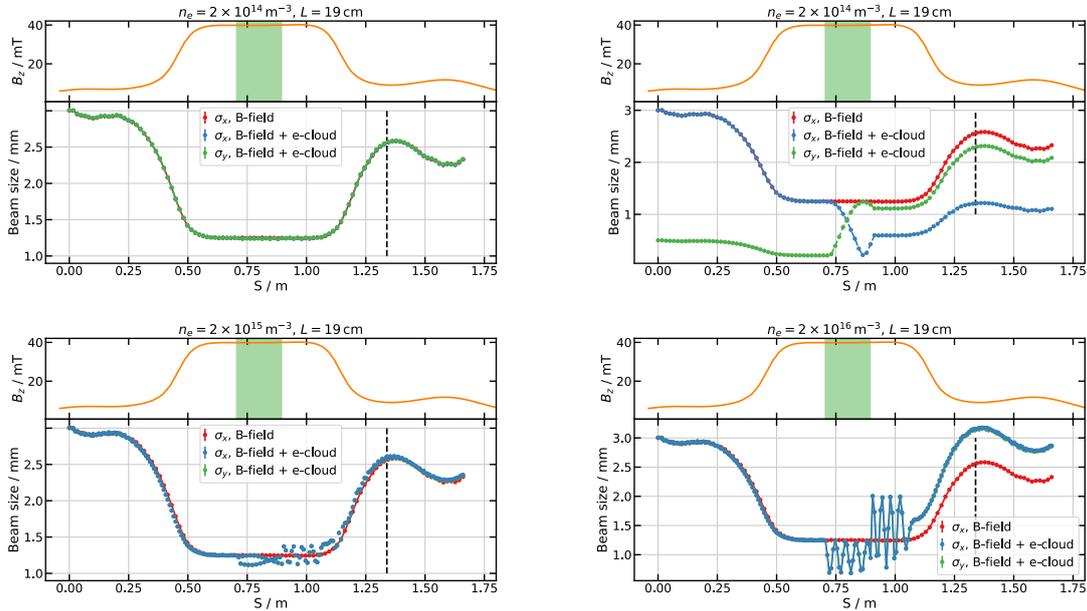
Figure E.5: The evolution of the transverse size ($\sigma_{x,y}$) of a beam of positrons that propagates along the beam-line when the storage trap is empty (B-field only) compared to plasma being confined (B-field + e-Cloud) with three electron densities of $2 \times 10^{14}\,\mathrm{m}^{-3}$ (top), $2 \times 10^{15}\,\mathrm{m}^{-3}$ (middle), and $2 \times 10^{16}\,\mathrm{m}^{-3}$ (bottom). An observable change in the beam size appears at the highest density. Due to the small separation of the planes at which the beam size is sampled, the faster change in the beam size due to the cyclotron rotation of the positrons can also be observed.

of it up to the position of the detector. For an electron density below $\sim 1 \times 10^{15}\,\mathrm{m}^{-3}$, no significant change in the beam size was observed. For higher density, $\sim 1 \times 10^{16}\,\mathrm{m}^{-3}$ and above, the beam size is larger at the position of the detector compared with the case of no plasma in the storage trap. The increase in the beam size at relatively high plasma densities is caused by the high amplitude oscillations of the beam size inside the electron cloud which perturbs the cyclotron motion of individual positrons downstream of the plasma. A mono-energetic beam passing through an ideal plasma lens has a well-defined focal point where the beam size is infinitesimal. Even though the beam energy spread is zero in the simulations, the beam size never reaches a value close to $0\,\mathrm{mm}$. This is an indication of the absence of focusing. This result suggests that, even at relatively high plasma densities, a focusing effect on the positron beam would not be observed. By comparison, in the absence of a magnetic field, a plasma with a density around $10^{14}\,\mathrm{m}^{-3}$ should focus a low energy positron beam in a space shorter than $1\,\mathrm{cm}$.

Furthermore, figure E.5 also shows several oscillations of the beam size inside the plasma which can be explained by the cyclotron motion of the positrons. The effect of the cyclotron motion of individual positrons on the beam size is visible due to the small separation between the planes at which the beam size was sampled. Due to the large cyclotron frequency and low beam energy, the positrons execute several gyrations during the transit time through the plasma.

Positron-beam rotation angle

The simulations presented above show an absence of an observable focusing effect of the electron plasma on a positron beam due to the presence of the external magnetic field of the trap in addition to the space-charge attraction of the plasma and due to the low energy of the positrons. Thus, a different experimental observable is needed to measure the impact of the plasma on the propagation of the beam. For example, an elliptical positron beam that propagates through the plasma on the central axis will rotate around this axis. The elliptical shape of the beam enables a measurement of the amount by which the beam rotates from the entry to the exit plane of the lens.

Alternatively. a thin pencil-beam that enters the electron plasma with a transverse offset from the central axis of the trap will undergo an azimuthal rotation around the axis and will exit the plasma trap at a different transverse location. It was argued previously that the positrons in the beam rotate with the same angular velocity which is determined by the plasma density and the solenoidal field. To verify this statement, an elliptical positron beam was tracked through the beam-line and its orientation in the xy plane was sampled at several illustrative locations. To sample the shape and orientation of the beam, the position of each positron in the transverse plane was recorded at several locations along the beam-line.

Figure E.6 shows how the orientation of a Gaussian elliptical beam that starts with $\sigma_x = 3\,\mathrm{mm}$, $\sigma_y = 0.5\,\mathrm{mm}$, and the major axis parallel to the x direction, changes through the beam-line. At the entry plane of the plasma, the beam has been focused due to the density of the magnetic field lines increasing upstream of the storage trap. Through the plasma, the beam rotates as a rigid object by almost two and a half full rotations.
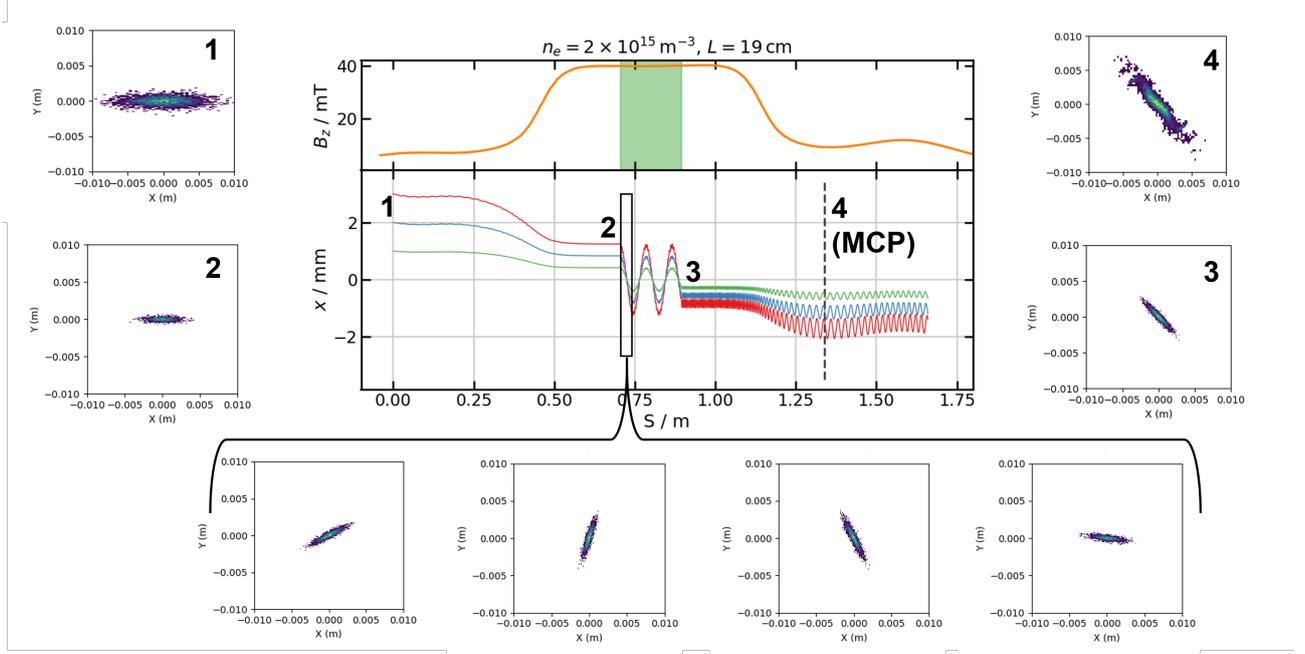
Figure E.6: Summary of single particle- and full beam-tracking results. The trajectory of individual positrons are shown with the red, blue, and green lines. The lateral panels show the orientation of an elliptical positron beam at various key locations in the beam-line: 1–start of the beam-line, 2–entry plane of the plasma, 3–exit plane of the plasma, 4–position of the detector. The four insets on the bottom show the rotation of the beam in the first approximately $4\,\mathrm{cm}$ of the plasma.

Between the exit plane of the plasma and the detector, the beam expands due to the decrease in the magnetic field strength. However, the rotation angle with which the beam exits the plasma is preserved up to the detector.

From the equations (E.1) and (E.2), one can obtain a new expression for the angle by which the beam rotates between the entry and exit planes of the plasma

$$\Delta\theta = \frac{e n_e L}{2\epsilon_0 B v_{\mathrm{beam}}} \tag{E.4}$$

where v_{beam} is the velocity of the positrons in the beam and L is the length of the plasma. Figure E.7 shows the amount by which the beam rotates within the electron plasma for three different positron energies, 50, 85, and $100\,\mathrm{eV}$, as a function of the figure of merit $n_e L$. The simple model of equation (E.4) agrees well with the beam-tracking simulations performed using BDSIM. The region of interest with rotation angles between about $0.02\,\mathrm{rad}$ and $pi\,\mathrm{rad}$ corresponds to $4\times10^{11}\,\mathrm{m^{-2}} \lesssim n_e L \lesssim 5\times10^{13}\,\mathrm{m^{-2}}$. In the treatment above, the beam was considered to be relatively large and to enter the beam-line on the beam axis. The same results apply for a thin pencil beam that starts to propagate with a radial offset. The final position of the pencil beam will be rotated by an equivalent angle also given by equation (E.4).

The beam-tracking simulation was also used to look at the correlation between the angle by which the beam rotates through the electron plasma and the rotation angle which can be measured at the position of the detector relative to the orientation of the beam at the start of the beam-line. Figure E.7 indicates that a good correlation exists only for relatively large angles above approximately $0.2\,\mathrm{rad}$. For lower angles, the beam rotation angle
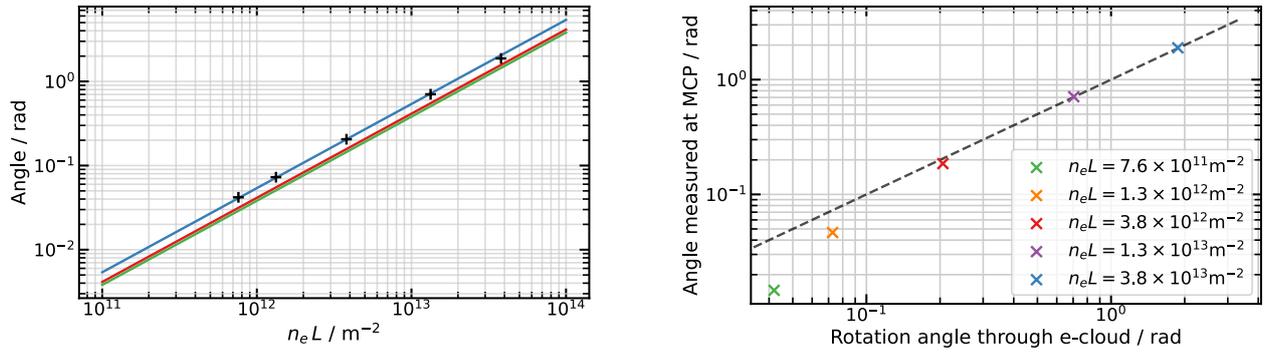
Figure E.7: **Left:** The angle by which a beam of positrons rotates while passing through an electron plasma of density n_e and length L for a beam energy of 50, 85, and 100 eV. The solid lines show the result of the simple model of equation (E.4). The crosses show the rotation of an elliptical positron beam through a uniform electron plasma according to beam-tracking simulations in BDSIM. **Right:** The correlation between the angle by which a positron beam rotates through the electron cloud and the rotation angle measured at the position of the detector (MCP) compared to the start of the beam-line. The dashed line marks the 1:1 correlation.

measured at the detector is less than the amount by which the beam rotates by passing through the plasma. This is a result of the rotation of the beam inside the plasma being smaller or comparable to the rotation caused by the propagation through the rest of the beam-line.

In conclusion, the propagation of a low-energy positron beam through an electron plasma confined within a Penning-Malmberg trap leads to an azimuthal rotation of the positrons around the central beam axis. Under the condition of a uniform electron plasma, the positrons rotate with a constant angular frequency which is independent of the radial position of the positrons as long as the beam propagates within the plasma. Given the layout of the positron beam-line at Swansea University, a measurable beam rotation could be obtained for plasma densities, n_e, in the range $4 \times 10^{11}\,\mathrm{m}^{-2} \lesssim n_e L \lesssim 5 \times 10^{13}\,\mathrm{m}^{-2}$, where L is the length of the plasma. In order for such experimental observations to be made, the radial size of the electron plasma has to be large enough to allow the positron beam to propagate with a measurable transverse offset from the central axis. Furthermore, the central region of the electron plasma needs to have a highly uniform charge density.

The particle-tracking studies outlined above establish the range of charge density of a confined electron plasma that is suitable for an experimental observation of the focusing effect on low-energy positron beams.
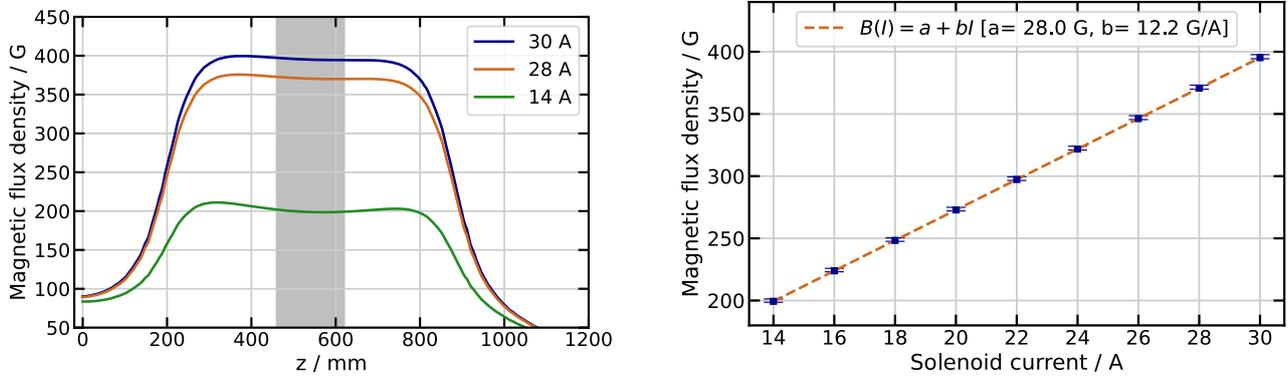
Figure E.8: Characterisation of the magnetic flux density along the positron beam-line at Swansea obtained with a model of the main solenoids and guiding coils in FEMM [223]. **Left:** The magnetic flux density along the axis of the solenoid of the electron trap for three values of the current through the magnet. The asymmetry is caused by the guiding coils upstream and downstream of the solenoid. The grey area indicate the region where the electrodes of the trap are positioned. **Right:** The axial component of the magnetic flux density at the centre of the trap as a function of the current through the solenoid.

Appendix F

The particle-in-cell (PIC) simulation method

F.1 PIC algorithm

Kinetic models are typically employed to investigate the dynamics of a non-neutral plasma when the time and length scales of interest are dictated by the plasma frequency and the Debye length, respectively. Several numerical techniques have been developed for kinetic simulations of a plasma. The studies included in this thesis are based on the particle-in-cell (PIC) method and on the VSim simulation software [233]. PIC codes simulate the dynamics of an ensemble of particles with the use of *macro-particles* which replace the continuous distribution function in phase-space by a discrete representation of it. Hence, macro-particles can be viewed as a cloud of many real particles or as an element of incompressible fluid moving in the 6-D phase-space. Each macro-particle has its own charge, q, and mass, m, depending on the number of real particles that it represents, but the ratio q/m is that of the real plasma species.

Macro-particles are also characterised by finite size. The PIC method is intended to simulate the collective behaviour of a plasma which is dictated by the simultaneous interaction of many particles through the slow fall-off part of the Coulomb force at large separation between the particles. However, the large Coulomb force between particles when they pass closely by each other gives rise to large impulses and collisional effects which are undesirable in a collisionless model. As a consequence, the rapidly varying forces associated with close encounters can be reduced with the use of finite-sized macro-particles [305]. One scheme to implement macro-particles with finite size is to smear out their charge on a finite region of space, typically on the cell of a grid that divides the simulation volume. Thus, variations in the charge density that are smaller than the size of a cell or a macro-particle cannot be resolved in a PIC simulation.

The use of macro-particles in VSim allows a simulation to be run in a reasonable amount of time without modelling all the physical particles while still capturing the kinetic effects from the particle distribution in space and time. VSim is based on the Vorpal computational engine which is built on the PIC algorithms of Hockney and Eastwood [306] and Birdsall and Langdon [307]. While there are over 70 variations of particle type and particle pusher available in VSim, the simulations presented in this chapter employed the 'relativistic Boris pusher' to update the particle velocities at each time step.

Boundary conditions at the edges of the simulation volume can be specified for the macro-particles. As a general rule, non-re-emitting absorbing surfaces were used at the faces of conducting materials and at the edges of the simulation volume. Particles incident on such a surface are completely removed from the simulation. Macro-particles are initialised at the beginning of the simulation by specifying their position, momentum, and the number of real particles that each macro-particle represents. Simultaneously, any external magnetic and electric fields are also defined at the initial time step of the simulation by specifying their components at each cell of the grid.

In summary, following the initialisation of the plasma macro-particles and the external fields, an electrostatic PIC simulation consists of a loop which is repeated at every time step. The main loop is composed of the following steps:

1. Updating the fields by integration of the fields equations, $\rho \to (\mathbf{E}, \mathbf{B})$

2. Interpolation of the forces at the position of the macro-particles, $(\mathbf{E}, \mathbf{B}) \to \mathbf{F_i}$

3. Integration of the equation of motion to update the particles position and velocity, $\mathbf{F_i} \to \mathbf{u_i} \to \mathbf{x_i}$

4. Interpolation to deposit the charge density at the location of the nodes of the grid, $(\mathbf{x_i}, \mathbf{u_i} \to \rho)$.

Moreover, a key aspect of the PIC method is the coupling of fields defined at discrete grid points with the continuous motion of the macro-particles.

Electric and magnetic fields

Important constitutive elements of the PIC simulation method are the electric and magnetic fieldsa as they dictate the motion of the charged particles. In general, the fields are represented as scalar or vector functions of space and time. In VSim, the fields are implemented as values on a grid, i.e. each cell has an associated value for a particular field. As the fields are most commonly part of differential equations that describe their evolution in time or space, the different components of a field have an associated location on a cell to enable the use of finite difference methods to evaluate the derivatives of the field. The location of a component of a field is the place where its value most accurately represents the real field value. Firstly, a scalar field, like the electrostatic potential, is defined at a single corner of each cell or at the cell centre. Secondly, for a vector field, the components are specified either at the centres of the edges or at the centres of the faces [308]. Some fields are used for deposition of charge and/or current in each cell.

Field updating

Fields in a PIC simulation can be either 'dynamic' or 'static', depending on whether the values of the fields attached to the simulation grid change between time steps or not. For example, the magnetic field in an electrostatic simulation is static. The values of the magnetic flux density are set once, maintained constant, and used throughout the simulation. On the other hand, in the case of both an electrostatic and an electromagnetic simulation, the electric field evolves at each time step. There are two procedures to update a field: explicit and implicit. The explicit update uses old field values to write the new field values at a given cell. The implicit

method requires an equation to be solved first in order to obtain the new field values. As most of the PIC simulations with VSim included in this chapter model static magnetic and time-dependent electric fields, an outline of the update procedure is given below for electrostatic fields.

Throughout all of the simulations presented in this chapter, the magnetic field was kept static. The values of the magnetic flux density were either imported at the beginning of the simulation from a field map generated by a separate code or defined from a single value under the assumption of a spatially uniform field. On the other hand, the electric field is found at each time step by solving for the potential, which satisfies Poisson's equation [308]

$$-\nabla \cdot \epsilon \left(\nabla \phi \right) = \rho, \tag{F.1}$$

and then calculating the electric field from

$$\mathbf{E} = -\nabla \phi. \tag{F.2}$$

The gradients are obtained via finite differences. The scalar potential is discretised by introducing the cell-averaged values ϕ_i. The discrete charge density is similarly defined as a set of average values, one for each cell. This procedure represents an implicit update since the potential, ϕ, has to be found first in order to update the electric field, $\mathbf{E}$. Therefore, the update starts by setting up a system of linear finite-difference equations to solve Poisson's equation. The system is solved by inverting a matrix at each time step and finding the unknown, ϕ, from the source term (charge density), ρ. The dielectric constant, ϵ, typically corresponds to that of vacuum, but it can be adapted to a different material or to a user-defined function over a specific region of space.

Upon discretisation in a 3-D space, Poisson's equation connects each node with six other adjacent ones and cannot be applied to the nodes at the edges of a simulation. On those grid nodes, one must apply boundary conditions (BC). Two types of BCs were used here; one specifies the value of the potential at a node (Dirichlet BC), the other specifies the value of the first derivative of the potential (Neumann BC). Furthermore, if the simulation includes volumes or surfaces of constant potential, all the nodes interior to that shape are fixed at the respective potential and the rest of the matrix is constructed as before.

Once the numerical matrix is built, finding the electric field at each time step requires a large matrix equation to be solved. These matrices become very large as the size of the problem increases since the number of elements in the matrix is close to the total number of cells in the simulation. VSim provides multiple direct or iterative solvers to cope with large matrices. A direct solver can be unsuitable for large problem sizes as the matrix can become too large to be held in memory. The alternative is to use an iterative method which finds the solution through successive matrix operations that converge to the solution. A choice of several pre-conditioners is available in VSim [244] to transform the linear system representing the Poisson equation into systems with more favourable convergence behaviour. Different combinations of solvers and pre-conditioners can be used in VSim. Typically in this work, several test runs were carried out before the actual full simulations to determine the optimal combination based on the CPU time required to solve Poisson's equation once at a single time step. It was found most often that the default settings provided by the user interface (the Visual Setup) resulted in the

fastest combination of solver type, solver parameters, and pre-conditioner.

It must be noted that the choice of using an electrostatic solver over a more complete electromagnetic one for the simulations included in this work is motivated by the size of the time step required in each case and, in turn, by the total CPU time required for a simulation to reach a certain total time. In the electrostatic (ES) case, the maximum size of the time step is restricted by the maximum thermal speed of the macro-particles and the minimum size of the cells in the grid. For an electromagnetic (EM) calculation, the maximum size of the time step is dictated by the speed of light and the minimum cell size. Thus, it was often the case that an EM solver required a time step three to four orders of magnitude smaller than the size of the time step suitable for an ES solver to limit numerical heating and avoid the unrealistic loss of macro-particles. As a result, the EM solver leads to the use of a significantly larger number of time steps to simulate a certain time interval compared to the ES solver. Even though the EM solver is faster in computing the fields at each time step, overall, it was found in this work that the ES solver is a better choice that leads to reduced requirements for total CPU time.

Besides the choice of the solver, i.e. how the fields are updated, several parameters of a PIC simulation must be chosen appropriately to ensure that the motion of the macro-particles is modelled correctly according to the physical laws that determine it. The most important parameters and the associated restrictions on their values are discussed below.

Choice of time step

To resolve the motion of a charged particle in a magnetic field properly, the time step should be smaller than the cyclotron period which is given by

$$\tau_c = \frac{2\pi m}{qB} \approx 3.57 \times 10^{-11} \frac{1}{B\,[\text{T}]} \quad [\text{s}]. \tag{F.3}$$

In general, the motion of a charged particle in a magnetic field can be treated as a superposition of a relatively fast circular motion around a point called the *guiding centre* and a relatively slow drift of this point. Typically, due to the high magnetic field and small gyroradius, the motion of a particle in a PIC simulation is close to a *guiding centre motion*, in which case, the small scale rotation of the particle around a magnetic field line is not fully resolved and the motion of the particle is a result of the slow drift of the guiding centre. Thus, larger time steps result in significant plasma heating.

Another natural time scale for the collective phenomena of interest in a plasma is the inverse plasma frequency ω_{pe}^{-1}, where for an electron plasma [309],

$$\omega_{pe} = 5.64 \times 10^4 \sqrt{n_e\,[\text{cm}^{-3}]} \quad [\text{rads}^{-1}]. \tag{F.4}$$

For a stable solution, $\Delta t \cdot \omega_{pe}^{-1} < 2$ is required. In the absence of other constraints, the use of a more restrictive time step, $\Delta t = 0.2\omega_{pe}^{-1}$ and down to $\Delta t = 0.05\omega_{pe}^{-1}$, can be used to give sufficiently accurate results [310]. The

above restriction on the time step is associated with the use of an explicit solver, a scheme that advances particles based on old forces from the previous time step. While explicit solvers are simpler and faster, an implicit scheme which uses a new field at the next time step to update particle velocity, requires a much smaller number of time steps.

A further constraint on the time step is set by the Courant–Friedrichs–Lewy (CFL) condition. For an electrostatic simulation, the effective CFL condition requires that a macro-particle moves at most one cell during a single time step. Hence, the presence of higher energy particles require smaller time steps. If a maximum temperature or a maximum thermal velocity, $v_{\mathrm{th,max}}$, is defined, the CFL condition can be satisfied and the plasma dynamics properly modelled by setting up the time step below the following minimum [244]

$$\min\left\{\tau_c, \frac{2}{\omega_{pe}}, \frac{1}{v_{th,max}\sqrt{1/\Delta x^2 + 1/\Delta y^2 + 1/\Delta z^2}}\right\}, \tag{F.5}$$

where $\Delta x, \Delta y, \Delta z$ are the cell dimensions along each of the x, y, z directions. The size of the time step is the defining factor that dictates the total running time required for a simulation to be completed on a fixed number of available CPUs.

Choice of grid cell size

In the PIC method, the short-range forces are only approximately modelled since their length scale is of the same order as the size of the mesh. This approximation and the discrete mesh lead to numerical heating and non-conservation of the total energy. However, the degree of numerical heating was observed to decrease by a comparable factor to the decrease in the cell size [311].

Simulations were carried out on a discrete Cartesian mesh with a transverse cell size $\Delta x = \Delta y$ and a longitudinal cell spacing Δz. The natural length scale in an electron plasma is given by the Debye length

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{n_e e^2}} \approx 7.43 \times 10^6 \sqrt{\frac{T_e\,[\mathrm{eV}]}{n_e\,[\mathrm{m}^{-3}]}} \quad [\mathrm{mm}], \tag{F.6}$$

which represents the upper limit for the mesh size in any direction. This condition adds to the difficulty of simulating a plasma with a low temperature T_e or a high density n_e.

Another length scale which is relevant for an electron moving in a magnetic field is the gyroradius (Larmor or cyclotron radius), ρ_c, defined as

$$\rho_c = \frac{m_e v_\perp}{eB} \approx 2.38 \times 10^{-3} \frac{\sqrt{T_e\,[\mathrm{eV}]}}{B\,[\mathrm{T}]} \quad [\mathrm{mm}], \tag{F.7}$$

where $v_\perp$ is the component of the velocity perpendicular to the magnetic field B. In a typical PIC simulation of a magnetised plasma, the scale of the gyroradius is smaller than the scale of the Debye length.

Electron collision rate

Non-neutral plasmas in a regime away from the Brillouin density limit satisfy the condition $\lambda_D > \rho_c$. In such a regime, when two electrons collide, their momentum vectors scatter. As a result, the guiding centre of each particle moves to a different position and leads to diffusion. Electron-electron Coulomb collisions occur at a rate given by the following formula [309]

$$\nu_{ee} = \frac{16}{15}\sqrt{\pi}\frac{n_e e^4}{T_e^{3/2} m_e^{1/2}} \ln\left(\frac{\rho_c T_e}{e^2}\right), \tag{F.8}$$

where cgs units are used and the temperature, T_e, has units of energy. Physically, ν_{ee} is the rate at which a thermal electron loses half of its momentum to other electrons [309]. Equation (F.8) is equivalent to

$$\nu_{ee} \approx 1.6 \times 10^{-6}\frac{n_e\,[\mathrm{cm}^{-3}]}{T_e^{3/2}\,[\mathrm{eV}^{3/2}]}\ln\left(6.9 \times 10^6 \rho_c\,[\mathrm{cm}]\,T_e\,[\mathrm{eV}]\right) \quad [\mathrm{Hz}]. \tag{F.9}$$

For example, a plasma with a density of the order $1 \times 10^9\,\mathrm{cm}^{-3}$, temperature $0.1\,\mathrm{eV}$ and confined by a field with $B = 30\,\mathrm{G}$ has an electron-electron collision rate $\nu_{ee} \approx 0.5\,\mathrm{MHz}$. This rate is a few orders of magnitude lower than the typical plasma frequency of the electron ensembles studied in the current chapter. Furthermore, the mean time between collisions is comparable to the magnitude of the time intervals for which electron plasmas were simulated in this work.

An ensemble of particles is considered to be weakly coupled when the thermal energy of the particles far exceeds the potential energy. In this regime, the trajectory of a particle at any time is affected by a large number of other particles. In other words, it is the mean field produced by the superposition of contributions from a large number of particles that dictates the motion of each particle in the plasma [225]. In general, there are a few conditions that determine whether a plasma is collisionless: large number of particles in a Debye cube with volume λ_D^3, kinetic energy much larger than potential energy, collision frequency much less than the plasma frequency, and a system with a size much larger than the mean free path of the particles.

Particle pusher

As described previously, once the electromagnetic fields have been computed at one time step of the PIC simulation, the macro-particles are advanced in time by updating their velocity and position. The integration of the respective equation of motion is done by the so-called 'particle pusher'. As the electric and magnetic fields are computed on fixed locations on the grid, they have firstly to be interpolated to the position of a particle.

The update of individual macro-particles is motivated by the observation that the time evolution of a distribution function in phase-space can be divided into two parts: advection in real space (movement) and advection in phase-space (acceleration). Typically, a leap-frog scheme is used in which the velocity and position are staggered in time: the particle position is advanced from time step n to $(n + 1)$, while the velocity is updated at half way

between time steps, from the step $(n - \frac{1}{2})$ to $(n + \frac{1}{2})$:

$$\frac{\mathbf{u_p}^{n+1/2} - \mathbf{u_p}^{n-1/2}}{\Delta t} = \frac{q_p}{m_p}\left(\mathbf{E}^n + \frac{1}{c}\frac{\mathbf{u_p}^{n+1/2} + \mathbf{u_p}^{n-1/2}}{2\gamma_p^n} \times \mathbf{B}^n\right), \tag{F.10}$$

$$\frac{\mathbf{x_p}^{n+1} - \mathbf{x_p}^n}{\Delta t} = \frac{\mathbf{u_p}^{n+1/2}}{\gamma^{n+1/2}}, \tag{F.11}$$

where $\gamma_p^n = \sqrt{1 + \left(\frac{\mathbf{u}_p^n}{c}\right)^2}$; $\mathbf{x_p}$ and $\mathbf{u_p}$ are the position and the four velocity of the macro-particle species p, respectively, and the t and Δt are the time and the time step in the simulation.

The standard algorithm for updating the particle velocity in PIC simulations is the Buneman-Boris scheme [307]. In this method, the velocity is updated through a series of steps that separate the contributions of the electric and magnetic field to the equation of motion and apply a combination of geometric rotations to update the velocity in a non-expensive explicit method [307].

Charge deposition and field interpolation

Updating the position and momentum of a particle requires information about the value of the electric and magnetic fields at the exact location of the macro-particle. In PIC simulations, the particles can be at any position inside the computation domain, the *continuum*, whereas the fields are defined on specific points attached to the grid of the simulation. Conversely, solving the field equations requires information about the charge density at the grid points.

The fields and the particles are, thus, linked by treating macro-particles as *super-particles* with finite extent. In this manner, macro-particles are characterised by the *shape function* which is symmetric, equates to zero outside a small range (compact support) and integrates to unity over the whole computation domain. At least first-order shape functions (linear interpolation), referred to as the *cloud-in-cell* method, are typically used by modern PIC codes. In order to ensure momentum conservation and avoid numerical instability, the shape function used to interpolate the fields to the location of the particles must be consistent with the scheme for charge deposition on the nodes of the grid.

For example, in the one-dimensional (in the x direction) case, the field, E_p, at the position, x_p, of a macro-particle is computed using the value of the field, $E(x_i)$, defined at the grid point i according to [225]

$$E_p(x_p) = \sum_i E(x_i)W(x_i - x_p) = \sum_i E_i W_i(x_p), \tag{F.12}$$

where x_i is the position of the grid point i. The formula above enables the reconstruction of the fields in the continuum using the assumption that the field is constant in each cell and equal to its cell-averaged value. W represents the *interpolation function* of the macro-particle, which is the influence of the position, x_p, of the macro-particle on the grid points close to it. The interpolation function is connected to the macro-particle shape

function, $S(x)$, through

$$W(x_i - x_p) = \int S_x(x - x_p)\Pi\left(\frac{x - x_i}{\Delta x}\right). \tag{F.13}$$

While $S(x)$ defines how the macro-particle charge is distributed about its centre (x_p), $W(x)$ is the convolution of $S(x)$ with the top hat function (Π) of span equal to the cell size (Δx). The usefulness of the interpolation functions is that they allow a direct computation of the average charge density over a cell, the components of the field in the continuum, and the equivalent force without the need for integration. In multi-dimensional simulations, the complete shape function is simply calculated as

$$S_{i,j,k}(\mathbf{x_p}) = S_i(x_p)S_j(y_p)S_k(z_p). \tag{F.14}$$

It must be noted that the PIC algorithm described above ensures conservation of the total momentum of the ensemble of particles. However, when momentum is conserved, the total energy is not. Energy and momentum conservation are mutually exclusive [312] in typical PIC schemes. Thus, it becomes important always to monitor the evolution of the total energy in a PIC simulation based on the methods presented here.

F.2   Initialisation of macro-particle phase-space

A discussion is included here on the importance of the initial state of the plasma in PIC simulations of space-charge (Gabor) lenses. The arguments are highlighted starting from an example of a simulation of the $m_\theta = 1$ diocotron rotation of a plasma column inside a lens.

Due to the influence of the transverse density profile of the plasma column on the diocotron mode frequency, the plasma has to be initialised in the simulation as close as possible to the actual stable state both in terms of the spatial density profile and the corresponding velocity distribution. Thus, one can avoid an initial transient phase in the simulation during which the plasma column changes from the initial configuration to a more stable state and also avoid numerical heating. Simply initialising the plasma with a uniform density has been observed in simulations to lead to an initial period of time required for the plasma to relax into a new stable state by typically expanding in the radial direction and contracting axially.

Figure F.1 shows an example from a simulation where the uniform initial plasma column relaxed into a significantly different longitudinal profile. In this example, the radial density distribution was calculated from the 1-D radial form of Poisson's equation (eq. (3.18)) assuming an infinitely long plasma column. A hard-edge uniform density was used in the longitudinal direction. The peak electron density was chosen below the Brillouin limit dictated by the magnitude of the magnetic field of the plasma trap, $3\,\text{mT}$. For the components of velocity at the start of the simulation, the plasma was assumed to undergo rigid rotation around its axis of symmetry. Thus, the radial component of the electron velocity was determined from the rigid-rotor angular velocity with which the radial density profile was calculated. Lastly, the axial component of the velocity was sampled from the Maxwell distribution with a finite temperature, $T_e = 1\,\text{eV}$. Such an initial state leads to a redistribution of
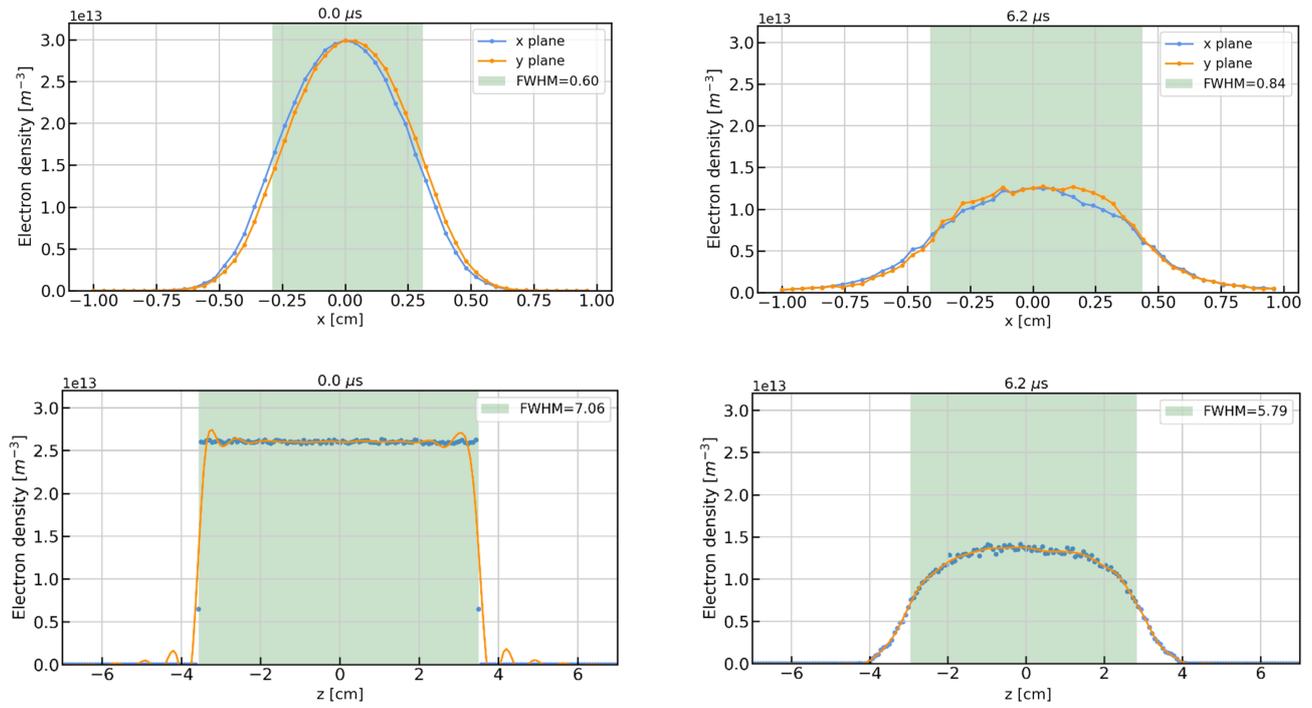
Figure F.1: The radial (top row) and longitudinal (bottom row) of an electron plasma column at the start of the simulation (left column) and at a later time step (right column). The initial plasma distribution in 6-D phase-space is set according to a finite-temperature global equilibrium and rigid rotation.

the electrons in the first part of the simulation, which, in turn, results in a significant drop in the peak plasma density and an increase of the radial FWHM of the plasma. Thus, it becomes more difficult to set the maximum plasma density and the transverse size of the column throughout a PIC simulation if the initial distribution of the plasma rapidly changes during the first part of a run.

Furthermore, the expansion of the plasma is accompanied by a more rapid loss of macro-particles compared to the rate of loss in the later stages of the simulation when the plasma reaches a stable state as shown in fig. F.2. The re-distribution of the electrons leads to heating of the plasma especially in the axial direction. Figure F.3 shows that the plasma temperature throughout the simulation is dictated by the initial phase of expansion. Even though the transverse and longitudinal temperatures are equal at the start of the simulation, they evolve without coupling given the short duration of the simulations described in this work compared to the inverse of the electron-electron collision frequency. The amount of plasma heating at the start of the simulation was slightly reduced by the use of a smaller time step.

One scheme to initialise the plasma closer to the stable state is to assume that, for small radial displacements of the column from the axis of the trap, the density distribution is not significantly different from the distribution of a plasma column in thermal equilibrium and with no offset from the axis of symmetry. The reasoning behind this method is that the equation which determines cylindrically symmetric non-neutral plasma equilibria of this kind have been extensively studied [50, 190, 194] and can be computed numerically. The numerical procedure has been constructed to compute the electrostatic equilibrium of a non-neutral plasma in a uniform magnetic field by solving Poisson's equation using an iterative method and is described in section 7.1.3. The class of equilibria described by the numerical calculation corresponds to a rigid rotation of the plasma about the axis of
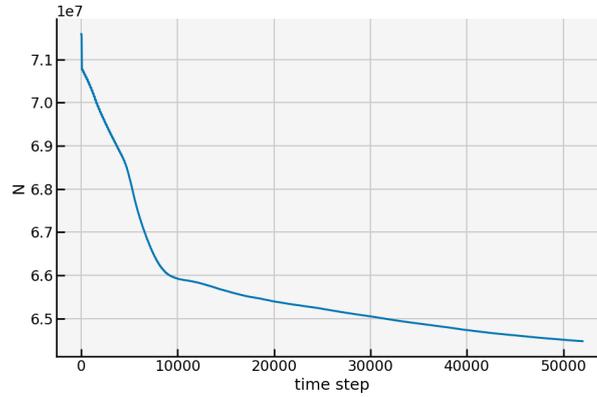
Figure F.2: Number of electron macro-particles lost from the plasma lens during the redistribution of the plasma shown in fig. F.1. The time step is $\Delta t = 90\,\text{ps}$.
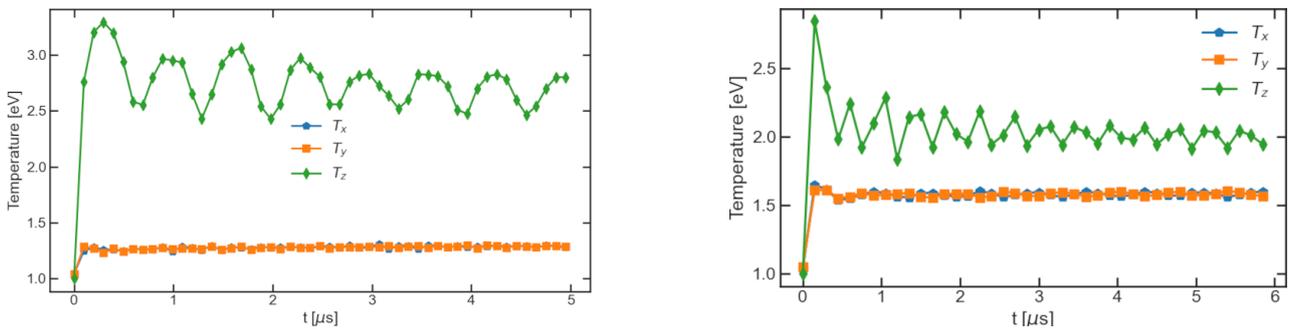
Figure F.3: The evolution of the transverse (x, y) and longitudinal (z) electron temperature in a PIC simulation with an initial transient phase of the plasma for two different time steps: $150\,\text{ps}$ (left) and $50\,\text{ps}$ (right).

symmetry with constant angular velocity, ω_{re}. In addition, the radial plasma density profile is flat up to the edge of the plasma where it quickly drops over a region with an extent of several Debye lengths (see fig. F.5).

The main advantage of setting the initial density distribution of the plasma in a PIC simulation of a space-charge lens according to the distribution of a plasma in global thermal equilibrium is that the spatial density profile does not have any hard edges. Thus, the peak electron density decreases smoothly at the radial edge of the plasma and at the two ends of the plasma column. As a result, there are only a few excess electrons near the edge of the plasma in addition to the electrons that form a stable plasma. These excess electrons are the ones that were observed to cause a transient phase at the start of the simulation as they are pushed back into the plasma by the confining fields and the self-potential of the plasma. In other words, creating fewer excess electrons at the start of the simulation decreases the source of subsequent plasma waves (travelling density perturbations) and changes in the overall shape and peak density of the plasma from the ones set initially. However, calculating the density distribution of a plasma in thermal equilibrium represents a step completely separate from the PIC simulation itself and adds to the overall computation time required to study a specific plasma behaviour.

Moreover, the stability of the a plasma state in the simulation is ensured by the distribution of electrons in the velocity space in addition to the spatial distribution of electron density. The Poisson-Boltzmann system of equations that describe a plasma in global thermal equilibrium provides a direct way of calculating the rigid-rotor angular velocity of the plasma once a self-consistent solution is found for the distribution of the plasma density. The radial component of the velocity of the electrons at the start of the simulation can be specified given the angular velocity of the plasma about its axis of symmetry. Lastly, the axial component of velocity is typically sampled from Maxwell's distribution with the temperature identical to the one used in the calculation of the density distribution. While the scheme described above to initialise the electrons in a PIC simulation of the Gabor lens leads to shorter transient stages in the first part of the simulation and less significant plasma re-distribution, these effects were still observable.

The solution found in this work for the plasma to be initialised in a stable state is based on a modification of the rigid-rotor assumption. In other words, the distribution of the electrons in the velocity space is calculated from the cold-fluid ($T_e = 0\,\mathrm{eV}$, no thermal pressure) theory of the plasma by assuming the spatial distribution of the density predicted by the theory of thermal equilibrium ($T_e > 0$). This approach allows the use of a soft-edge, continuous plasma profile in the PIC simulations with the associated velocity distribution that would maintain the initial density distribution of electrons stable throughout the entire simulation.

The reason for choosing the 'hybrid' scheme outlined in the previous paragraph to initialise the electrons in a PIC simulation is that VSim, the code used in this work, is a collision-less kinetic code which does not model the pressure of an ensemble of particles when macro-particles are employed. Thus, an electron-plasma column that rotates about its axis of symmetry with a given density distribution will remain stable during a PIC simulation if the centrifugal repulsion is balanced by the Lorentz force due to the self-potential of the plasma
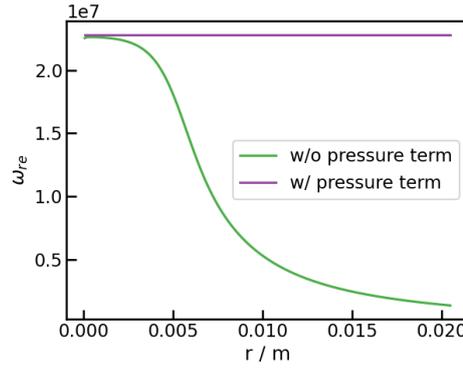
Figure F.4: Angular velocity of a fluid element as a function of radius for an infinitely long electron plasma column ($n_e = 1 \times 10^{14}\,\mathrm{m^{-3}}$) in equilibrium inside a Penning-Malmberg trap ($40\,\mathrm{mT}$ axial magnetic flux density). In the absence of elastic collisions and, thus, pressure, the plasma column does not rotate as a rigid rotor.

and the external fields of the trap. The radial component of the equilibrium condition can be written as [187]

$$\omega_{re} = \frac{E_r}{Br} + \frac{k_B T_e}{eBrn_e}\frac{\partial n_e}{\partial r} + \frac{\omega_{re}^2}{\omega_{ce}}, \tag{F.15}$$

where $E_r = -\partial\phi/\partial r$ is the radial component of electric field, and ω_{ce} is the electron cyclotron frequency. Equation (F.15) shows that a plasma column with zero temperature and a density that decreases smoothly with the radius requries a radially-dependent angular velocity for radial equilibrium to be achieved. Figure F.4 shows the solution of eq. (F.15) both in the presence and absence of the pressure term for the initial radial density profile given in fig. F.5. For a cold-fluid plasma the angular velocity needs to decrease radially towards the edge of the plasma column in order to maintain a pseudo-flat radial density profile. In other words, in the absence of pressure, the electrons near the edge of the plasma must rotate with lower azimuthal velocity since the $\mathbf{v} \times \mathbf{B}$ contribution to the Lorentz force has to compensate for a smaller force pushing the electrons radially outwards due to the interaction with the rest of the plasma.

In summary, a 'hybrid' scheme was devised in this work to initialise the electrons in a PIC simulation such that the plasma remains stable after the start of the simulation. The spatial distribution of electron density is calculated numerically from the finite temperature theory of the plasma. Subsequently, the angular velocity of the electrons is calculated from eq. F.15 using the radial gradient of the plasma density obtained in the previous step. For the axial component of the electron velocity, values are sampled from Maxwell's distribution with the same electron temperature as the one used in the calculation of the distribution of the electron density. Figures F.5 and F.6 show the changes in the radial and longitudinal profiles of the plasma density from the start to the end of a PIC simulation of a plasma column undergoing diocotron rotation around the central axis of a Penning-Malmberg trap. The variations of the length and the radial FWHM of the plasma are both below approximately 3%.
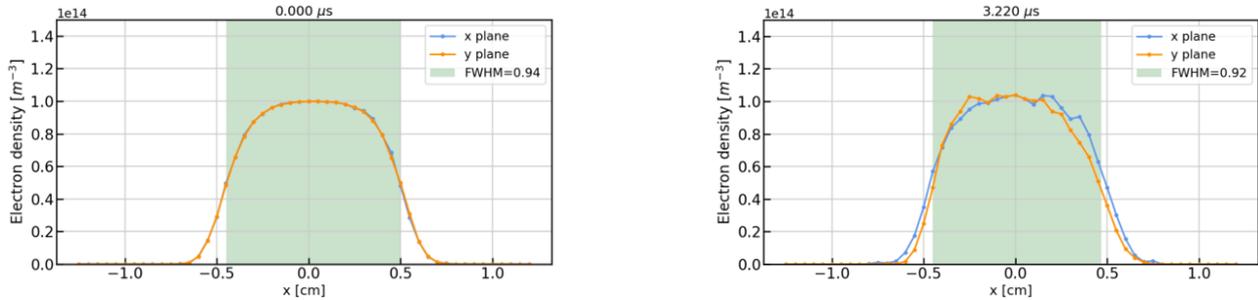
Figure F.5: The electron density in the mid-lane of the lens as a function of the transverse coordinates x and y before the start of the PIC simulation (left) and at a later time step after approximately one diocotron rotation (right). The shaded area marks the region where the plasma density is greater than half of its maximum value.
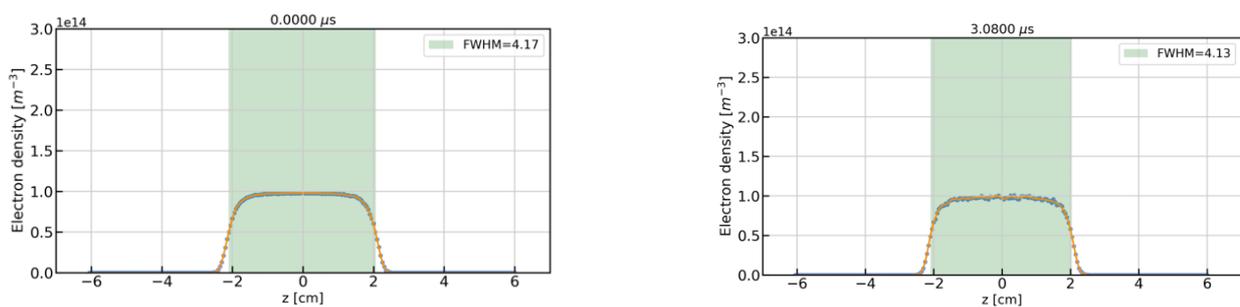
Figure F.6: The electron density on the central axis of the lens before the start of the PIC simulation (left) and at a later time step after approximately one diocotron rotation (right). The blue markers represent the data. A smoothing line is added in orange to guide the eye. The shaded area marks the region where the plasma density is greater than half of its maximum value.

Bibliography

[1] The World Health Organisation. Cancer. `https://www.who.int/news-room/fact-sheets/detail/cancer`, 2023. (see p. 27)

[2] Worldwide cancer statistics | Cancer Research UK `https://www.cancerresearchuk.org/health-professional/cancer-statistics/worldwide-cancer`. Accessed: 2023-05-01 (see p. 27)

[3] N. R. Datta, M. Samiei, and S. Bodis. Radiation Therapy Infrastructure and Human Resources in Low- and Middle-Income Countries: Present Status and Projections for 2020. *International Journal of Radiation Oncology, Biology, Physics*, 89: 448–457, 2014. DOI: `10.1016/j.ijrobp.2014.03.002` (see pp. 27, 28)

[4] H. H. Chen and M. Tien Kuo. Improving radiotherapy in cancer treatment: Promises and challenges. *Oncotarget*, 8: 62742–62758, 2017. DOI: `10.18632/oncotarget.18409` (see p. 27)

[5] J. M. Borras *et al.* How many new cancer patients in Europe will require radiotherapy by 2025? An ESTRO-HERO analysis. *Radiotherapy and Oncology*, 119: 5–11, 2016. DOI: `10.1016/j.radonc.2016.02.016` (see p. 27)

[6] R. Atun *et al.* Expanding global access to radiotherapy. *The Lancet Oncology*, 16: 1153 –1186, 2015. DOI: `10.1016/S1470-2045(15)00222-3` (see p. 27)

[7] T. E. Merchant, H. M. Conklin, S. Wu, R. H. Lustig, and X. Xiong. Late Effects of Conformal Radiation Therapy for Pediatric Patients With Low-Grade Glioma: Prospective Evaluation of Cognitive, Endocrine, and Hearing Deficits. *Journal of Clinical Oncology*, 27: PMID: 19581535, 3691–3697, 2009. DOI: `10.1200/JCO.2008.21.2738` (see p. 27)

[8] M. Baumann *et al.* Radiation oncology in the era of precision medicine. *Nature Reviews Cancer 2016 16:4*, 16: 234–249, 2016. DOI: `10.1038/nrc.2016.18` (see p. 27)

[9] G. Aymar *et al.* LhARA: The Laser-hybrid Accelerator for Radiobiological Applications. *Frontiers in Physics*, 8: 2020. DOI: `10.3389/fphy.2020.567738` (see pp. 27, 29, 30, 52, 64–69, 155, 187, 192, 195)

[10] P. Allport *et al.* The Laser-hybrid Accelerator for Radiobiological Applications: R&D proposal for the preliminary, pre-construction phases. CCAP-TN-10 `https://ccap.hep.ph.ic.ac.uk/trac/raw-attachment/wiki/Communication/Notes/CCAP-TN-10.pdf` The LhARA Collaboration, 2022. (see pp. 27, 30, 64, 66, 153)

[11] M.-C. Vozenin *et al.* The advantage of Flash radiotherapy confirmed in mini-pig and cat-cancer patients. *Clinical Cancer Research*, 25: 35–42, 2019. DOI: `10.1158/1078-0432.CCR-17-3375` (see pp. 28, 62)

[12] V. Favaudon *et al.* Ultrahigh dose-rate FLASH irradiation increases the differential response between normal and tumor tissue in mice. *Science Translational Medicine*, 6: 245ra93 LP –245ra93, 2014. DOI: `10.1126/scitranslmed.3008973` (see pp. 28, 62)

[13] J. R. Hughes and J. L. Parsons. FLASH Radiotherapy: Current Knowledge and Future Insights Using Proton-Beam Therapy. *International Journal of Molecular Sciences*, 21: 2020. DOI: `10.3390/ijms21186492` (see pp. 28, 63)

[14] Y. Prezado *et al.* Proton minibeam radiation therapy spares normal rat brain: Long-Term Clinical, Radiological and Histopathological Analysis. *Scientific Reports*, 7: 14403, 2017. DOI: `10.1038/s41598-017-14786-y` (see pp. 28, 63)

[15] Y. Prezado *et al.* Tumor Control in RG2 Glioma-Bearing Rats: A Comparison Between Proton Minibeam Therapy and Standard Proton Therapy. *International Journal of Radiation Oncology, Biology, Physics*, 104: 266–271, 2019. DOI: `10.1016/j.ijrobp.2019.01.080` (see pp. 28, 63)

[16] Particle Therapy Co-Operative Group `https://www.ptcog.ch`. Accessed: 2022-01-13 (see p. 28)

[17] Particle Therapy Co-Operative Group; Facilities in operation (October 2022 update) `https://www.ptcog.ch/index.php/facilities-in-operation`. Accessed: 2022-11-22 (see pp. 28, 59, 60)

[18] R. Mohan and D. Grosshans. Proton therapy – Present and future. *Advanced Drug Delivery Reviews*, 109: Radiotherapy for cancer: present and future, 26–44, 2017. DOI: `10.1016/j.addr.2016.11.006` (see pp. 28, 61)

[19] S. V. Bulanov and V. S. Khoroshkov. Feasibility of using laser ion accelerators in proton therapy. *Plasma Physics Reports*, 28: 453–456, 2002. DOI: `10.1134/1.1478534` (see pp. 28, 42)

[20] E. Fourkal *et al.* Particle in cell simulation of laser-accelerated proton beams for radiation therapy. *Medical Physics*, 29: 2788–2798, 2002. DOI: `10.1118/1.1521122` (see pp. 28, 42)

[21] V. Malka *et al.* Practicability of protontherapy using compact laser systems. *Medical Physics*, 31: 1587–1592, 2004. DOI: `10.1118/1.1747751` (see pp. 28, 42)

[22] H. Daido, M. Nishiuchi, and A. S. Pirozhkov. Review of laser-driven ion sources and their applications. *Reports on Progress in Physics*, 75: 056401, 2012. DOI: `10.1088/0034-4885/75/5/056401` (see pp. 28, 33, 34, 37, 40, 44)

[23] P. Chaudhary *et al.* Radiobiology Experiments With Ultra-high Dose Rate Laser-Driven Protons: Methodology and State-of-the-Art. *Frontiers in Physics*, 9: 2021. DOI: `10.3389/fphy.2021.624963` (see pp. 28, 40, 42, 63)

[24] G. A. P. Cirrone *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*. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. may 2013. 87791I DOI: `10.1117/12.2026530` (see p. 28)

[25] J. Bin *et al.* A new platform for ultra-high dose rate radiobiological research using the BELLA PW laser proton beamline. *Scientific Reports 2022 12:1*, 12: 1–15, 2022. DOI: `10.1038/s41598-022-05181-3` (see pp. 28, 29, 41, 46, 51)

[26] F. E. Brack *et al.* Spectral and spatial shaping of laser-driven proton beams using a pulsed high-field magnet beamline. *Scientific Reports 2020 10:1*, 10: 1–12, 2020. DOI: `10.1038/s41598-020-65775-7` (see pp. 28, 29, 46, 49)

[27] F. Kroll *et al.* Tumour irradiation in mice with a laser-accelerated proton beam. *Nature Physics*, 18: 316–322, 2022. DOI: `10.1038/s41567-022-01520-3` (see pp. 28, 42)

[28] J. G. Zhu *et al.* Experimental demonstration of a laser proton accelerator with accurate beam control through image-relaying transport. *Phys. Rev. Accel. Beams*, 22: 061302, 2019. DOI: `10.1103/PhysRevAccelBeams.22.061302` (see pp. 29, 46, 48)

[29] J. G. Zhu *et al.* Demonstration of tailored energy deposition in a laser proton accelerator. *Phys. Rev. Accel. Beams*, 23: 121304, 2020. DOI: `10.1103/PhysRevAccelBeams.23.121304` (see pp. 29, 46)

[30] M. Schollmeier *et al.* Controlled Transport and Focusing of Laser-Accelerated Protons with Miniature Magnetic Devices. *Phys. Rev. Lett.*, 101: 055004, 2008. DOI: `10.1103/PhysRevLett.101.055004` (see pp. 29, 46, 48)

[31] M. Nishiuchi *et al.* Focusing and spectral enhancement of a repetition-rated, laser-driven, divergent multi-MeV proton beam using permanent quadrupole magnets. *Applied Physics Letters*, 94: 061107, 2009. DOI: `10.1063/1.3078291` (see pp. 29, 46, 48)

[32] F. Schillaci *et al.* Characterization of the ELIMED Permanent Magnets Quadrupole system prototype with laser-driven proton beams. *Journal of Instrumentation*, 11: T07005, 2016. DOI: `10.1088/1748-0221/11/07/T07005` (see pp. 29, 46, 48)

[33]   K. Harres *et al.* Beam collimation and transport of quasineutral laser-accelerated protons by a solenoid field. *Physics of Plasmas*, 17: 023107, 2010. DOI: 10.1063/1.3299391 (see pp. 29, 46, 49)

[34]   T. Burris-Mog *et al.* Laser accelerated protons captured and transported by a pulse power solenoid. *Phys. Rev. ST Accel. Beams*, 14: 121301, 2011. DOI: 10.1103/PhysRevSTAB.14.121301 (see pp. 29, 46)

[35]   S. Busold *et al.* Commissioning of a compact laser-based proton beam line for high intensity bunches around 10 MeV. *Phys. Rev. ST Accel. Beams*, 17: 031302, 2014. DOI: 10.1103/PhysRevSTAB.17.031302 (see pp. 29, 46)

[36]   K. Zeil *et al.* Dose-controlled irradiation of cancer cells with laser-accelerated proton pulses. *Applied Physics B 2012 110:4*, 110: 437–444, 2012. DOI: 10.1007/S00340-012-5275-3 (see pp. 29, 49)

[37]   F. Dothan, H. Riege, E. Boggasch, and K. Frank. Dynamics of a z pinch for focusing high-energy charged particles. *Journal of Applied Physics*, 62: 3585–3591, 1987. DOI: 10.1063/1.339286 (see pp. 29, 50)

[38]   A. Tauschwitz *et al.* The plasma lens solution for heavy-ion beam focusing. *Il Nuovo Cimento A (1965-1970) 1993 106:11*, 106: 1733–1737, 2008. DOI: 10.1007/BF02821274 (see pp. 29, 50)

[39]   M. Turner *et al.* Radial density profile and stability of capillary discharge plasma waveguides of lengths up to 40 cm. *High Power Laser Science and Engineering*, 9: e17, 2021. DOI: 10.1017/hpl.2021.6 (see pp. 29, 51, 204)

[40]   R. M. Mobley, G. Gammel, and A. W. Maschke. Gabor Lenses. *IEEE Transactions on Nuclear Science*, 26: 3112–3114, 1979. DOI: 10.1109/TNS.1979.4329957 (see pp. 29, 84)

[41]   R. Booth and H. Lefevre. Space charge lens for high current ion beams. *Nuclear Instruments and Methods*, 151: 143–147, 1978. DOI: 10.1016/0029-554X(78)90481-0 (see pp. 29, 84)

[42]   H. Lefevre, R. Connolly, G. Sieger, and J. Overley. Scanning MeV-ion microprobe for light and heavy ions. *Nuclear Instruments and Methods in Physics Research*, 218: 39–42, 1983. DOI: 10.1016/0167-5087(83)90951-1 (see pp. 29, 84, 89)

[43]   J. A. Palkovic *et al.* Measurements on a Gabor Lens for Neutralizing and Focusing a 30-KeV Proton Beam. FERMILAB-CONF-88-177, 88-10-03 october 1988 (see pp. 29, 30, 84)

[44]   J. Pozimski, R. Dölling, P. Gross, and T. Weis. First Experimental Studies of a Gabor-Plasma-Lens in Frankfurt. 1992. (see pp. 29, 30, 85)

[45]   A. Goncharov, I. Protsenko, G. Yushkov, and I. G. Brown. Development of a high-current plasma lens for focusing broad beams of heavy metal ions. *Review of Scientific Instruments*, 71: 1401–1404, 2000. DOI: 10.1063/1.1150469 (see pp. 29, 207)

[46] J. Pozimski and O. Meusel. Space charge lenses for particle beams. *Review of Scientific Instruments*, 76: 063308, 2005. DOI: `10.1063/1.1904203` (see pp. 29, 30, 82, 86, 87, 89)

[47] K. Schulte *et al.* "Gabor Lens Performance Studies at the GSI High Current Test Injector" in *Proceedings of the 4th International Particle Accelerator Conference (IPAC 2013)*: (Shanghai, China). volume 7 JACoW Publishing, Geneva, Switzerland, 2013. TUPWO008, 3806–3808 (see pp. 29, 87, 88, 235)

[48] K. I. Thoma *et al.* Experiments with stable confined electron columns. 2020. DOI: `10.23732/CYRCP-2020-007.143` (see pp. 29, 86)

[49] D. Gabor. A Space-Charge Lens for the Focusing of Ion Beams. *Nature*, 160: 89–90, 1947. DOI: `10.1038/160089b0` (see pp. 29, 50, 68, 79, 83, 93, 214)

[50] J. H. Malmberg and J. S. deGrassie. Properties of Nonneutral Plasma. *Phys. Rev. Lett.*, 35: 577–580, 1975. DOI: `10.1103/PhysRevLett.35.577` (see pp. 29, 259)

[51] K. Schulte. *Studies on the focusing performance of a Gabor lens depending on nonneutral plasma properties*. PhD Thesis. Johann Wolfgang Goethe-Universität in Frankfurt am Main, 2014. (see pp. 29, 30, 82, 87–89)

[52] Y. Qingchang and Q. Hong. Self-sustaining space-charge lens. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 311: 14–18, 1992. DOI: `10.1016/0168-9002(92)90845-U` (see pp. 29, 87, 89, 210)

[53] O. Meusel, M. Droba, B. Glaeser, and K. Schulte. Experimental studies of stable confined electron clouds using Gabor lenses. *AIP Conf. Proc. C*, 1206051: Comments: 4 pages, contribution to the Joint INFN-CERN-EuCARD-AccNet Workshop on Electron-Cloud Effects: ECLOUD'12; 5-9 Jun 2012, La Biodola, Isola d'Elba, Italy, 157–160, 2013. DOI: `10.5170/CERN-2013-002.157` (see pp. 29, 86, 87, 210)

[54] T. Nonnenmacher *et al.* Anomalous Beam Transport through Gabor (Plasma) Lens Prototype. *Applied Sciences*, 11: 2021. DOI: `10.3390/app11104357` (see pp. 30, 92, 93, 98, 99, 111, 113, 117, 201)

[55] E. L. Clark *et al.* Measurements of Energetic Proton Transport through Magnetized Plasma from Intense Laser Interactions with Solids. *Phys. Rev. Lett.*, 84: 670–673, 2000. DOI: `10.1103/PhysRevLett.84.670` (see p. 33)

[56] A. Maksimchuk, S. Gu, K. Flippo, D. Umstadter, and V. Y. Bychenkov. Forward Ion Acceleration in Thin Films Driven by a High-Intensity Laser. *Phys. Rev. Lett.*, 84: 4108–4111, 2000. DOI: `10.1103/PhysRevLett.84.4108` (see p. 33)

[57] R. A. Snavely *et al.* Intense High-Energy Proton Beams from Petawatt-Laser Irradiation of Solids. *Phys. Rev. Lett.*, 85: 2945–2948, 2000. DOI: `10.1103/PhysRevLett.85.2945` (see pp. 33, 36)

[58] S. C. Wilks *et al.* Energetic proton generation in ultra-intense laser–solid interactions. *Physics of Plasmas*, 8: 542–549, 2001. DOI: `10.1063/1.1333697` (see p. 33)

[59]   A. P. L. Robinson, M Zepf, S Kar, R. G. Evans, and C Bellei. Radiation pressure acceleration of thin foils with circularly polarized laser pulses. *New Journal of Physics*, 10: 013021, 2008. DOI: `10.1088/1367-2630/10/1/013021` (see p. 33)

[60]   L. Yin *et al.* Monoenergetic and GeV ion acceleration from the laser breakout afterburner using ultrathin targets. *Physics of Plasmas*, 14: 056706, 2007. DOI: `10.1063/1.2436857` (see p. 33)

[61]   J. Fernández *et al.* Fast ignition with laser-driven proton and ion beams. *Nuclear Fusion*, 54: 054006, 2014. DOI: `10.1088/0029-5515/54/5/054006` (see p. 34)

[62]   A. Macchi, M. Borghesi, and M. Passoni. Ion acceleration by superintense laser-plasma interaction. *Rev. Mod. Phys.*, 85: 751–793, 2013. DOI: `10.1103/RevModPhys.85.751` (see pp. 33, 44)

[63]   L. A. Gizzi. "Laser-Driven Sources of High Energy Particles and Radiation" in*Laser-Driven Sources of High Energy Particles and Radiation*. Springer, 2019. 1–24 (see pp. 34, 35)

[64]   A. J. Mackinnon *et al.* Enhancement of Proton Acceleration by Hot-Electron Recirculation in Thin Foils Irradiated by Ultraintense Laser Pulses. *Phys. Rev. Lett.*, 88: 215006, 2002. DOI: `10.1103/PhysRevLett.88.215006` (see pp. 35, 37)

[65]   M. Roth and M. Schollmeier. Ion Acceleration—Target Normal Sheath Acceleration. *CERN Yellow Reports*, Vol 1 (2016): Proceedings of the 2014 CAS–CERN Accelerator School: Plasma Wake Acceleration, 2016. DOI: `10.5170/CERN-2016-001.231` (see pp. 36, 38)

[66]   J. Pozimski, M. Aslaninejad, and P. Posocco. "Advanced Gabor Lens Lattice for Laser Driven Hadron Therapy and Other Applications" in *Proc. of International Particle Accelerator Conference (IPAC'16), Busan, Korea, May 8-13, 2016*: (Busan, Korea). International Particle Accelerator Conference 7 Geneva, Switzerland: JACoW, june 2016. 1595–1597 DOI: `10.18429/JACoW-IPAC2016-TUPMY023` (see pp. 36, 41, 68, 91)

[67]   P. Mora. Plasma Expansion into a Vacuum. *Phys. Rev. Lett.*, 90: 185002, 2003. DOI: `10.1103/PhysRevLett.90.185002` (see pp. 36, 37)

[68]   P. Mora. Thin-foil expansion into a vacuum. *Phys. Rev. E*, 72: 056401, 2005. DOI: `10.1103/PhysRevE.72.056401` (see pp. 36, 37)

[69]   P. McKenna *et al.* Characterization of proton and heavier ion acceleration in ultrahigh-intensity laser interactions with heated target foils. *Phys. Rev. E*, 70: 036405, 2004. DOI: `10.1103/PhysRevE.70.036405` (see p. 36)

[70]   M. Nishiuchi *et al.* The laser proton acceleration in the strong charge separation regime. *Physics Letters A*, 357: 339–344, 2006. DOI: `10.1016/j.physleta.2006.04.053` (see p. 36)

[71]   D. C. Carroll *et al.* Carbon ion acceleration from thin foil targets irradiated by ultrahigh-contrast, ultraintense laser pulses. *New Journal of Physics*, 12: 045020, 2010. DOI: `10.1088/1367-2630/12/4/045020` (see p. 36)

[72] J. Fuchs *et al.* Spatial Uniformity of Laser-Accelerated Ultrahigh-Current MeV Electron Propagation in Metals and Insulators. *Phys. Rev. Lett.*, 91: 255002, 2003. DOI: `10.1103/PhysRevLett.91.255002` (see p. 37)

[73] Y. Sentoku *et al.* High-energy ion generation in interaction. of short laser pulse with high-density plasma. *Applied Physics B 2002 74:3*, 74: 207–215, 2014. DOI: `10.1007/S003400200796` (see p. 37)

[74] F. Wagner *et al.* Maximum Proton Energy above 85 MeV from the Relativistic Interaction of Laser Pulses with Micrometer Thick CH_2 Targets. *Phys. Rev. Lett.*, 116: 205002, 2016. DOI: `10.1103/PhysRevLett.116.205002` (see p. 37)

[75] A. Higginson *et al.* Near-100 MeV protons via a laser-driven transparency-enhanced hybrid acceleration scheme. *Nature Communications 2018 9:1*, 9: 1–9, 2018. DOI: `10.1038/s41467-018-03063-9` (see p. 37)

[76] E. L. Clark *et al.* Energetic Heavy-Ion and Proton Generation from Ultraintense Laser-Plasma Interactions with Solids. *Phys. Rev. Lett.*, 85: 1654–1657, 2000. DOI: `10.1103/PhysRevLett.85.1654` (see p. 37)

[77] W. J. Ma *et al.* Laser Acceleration of Highly Energetic Carbon Ions Using a Double-Layer Target Composed of Slightly Underdense Plasma and Ultrathin Foil. *Phys. Rev. Lett.*, 122: 014803, 2019. DOI: `10.1103/PhysRevLett.122.014803` (see p. 37)

[78] A. McIlvenny *et al.* Selective Ion Acceleration by Intense Radiation Pressure. *Phys. Rev. Lett.*, 127: 194801, 2021. DOI: `10.1103/PhysRevLett.127.194801` (see p. 37)

[79] B. M. Hegelich *et al.* Spectral properties of laser-accelerated mid-Z MeV/u ion beams. *Physics of Plasmas*, 12: 056314, 2005. DOI: `10.1063/1.1915350` (see p. 37)

[80] B. M. Hegelich *et al.* Laser acceleration of quasi-monoenergetic MeV ion beams. *Nature*, 439: 441–444, 2006. DOI: `10.1038/nature04400` (see p. 37)

[81] K. Kondo *et al.* High-Intensity Laser-Driven Oxygen Source from CW Laser-Heated Titanium Tape Targets. *Crystals*, 10: 2020. DOI: `10.3390/cryst10090837` (see p. 37)

[82] D. Jung *et al.* Laser-driven 1 GeV carbon ions from preheated diamond targets in the break-out afterburner regime. *Physics of Plasmas*, 20: 083103, 2013. DOI: `10.1063/1.4817287` (see p. 37)

[83] F. Nürnberg *et al.* Radiochromic film imaging spectroscopy of laser-accelerated proton beams. *Review of Scientific Instruments*, 80: 033301, 2009. DOI: `10.1063/1.3086424` (see p. 37)

[84] N. Chauvin. Space-Charge Effect. 2014. DOI: `10.5170/CERN-2013-007.63` (see p. 38)

[85] M. Borghesi *et al.* Electric field detection in laser-plasma interaction experiments via the proton imaging technique. *Physics of Plasmas*, 9: 2214–2220, 2002. DOI: `10.1063/1.1459457` (see p. 40)

[86]   W. Bang *et al.* Uniform heating of materials into the warm dense matter regime with laser-driven quasimonoenergetic ion beams. *Phys. Rev. E*, 92: 063101, 2015. DOI: `10.1103/PhysRevE.92.063101` (see p. 40)

[87]   M. Roth *et al.* Bright Laser-Driven Neutron Source Based on the Relativistic Transparency of Solids. *Phys. Rev. Lett.*, 110: 044802, 2013. DOI: `10.1103/PhysRevLett.110.044802` (see p. 40)

[88]   S. Fritzler *et al.* Proton beams generated with high-intensity lasers: Applications to medical isotope production. *Applied Physics Letters*, 83: 3039–3041, 2003. DOI: `10.1063/1.1616661` (see p. 40)

[89]   L. Torrisi *et al.* Ion implantation induced by Cu ablation at high laser fluence. *Applied Surface Science*, 252: 8533–8538, 2006. DOI: `10.1016/j.apsusc.2005.11.071` (see p. 40)

[90]   M. Roth *et al.* Fast Ignition by Intense Laser-Accelerated Proton Beams. *Phys. Rev. Lett.*, 86: 436–439, 2001. DOI: `10.1103/PhysRevLett.86.436` (see pp. 40, 43)

[91]   K. W. D. Ledingham, P. R. Bolton, N. Shikazono, and C.-M. C. Ma. Towards Laser Driven Hadron Cancer Radiotherapy: A Review of Progress. *Applied Sciences*, 4: 402–443, 2014. DOI: `10.3390/app4030402` (see p. 40)

[92]   N. Dover *et al.* Demonstration of repetitive energetic proton generation by ultra-intense laser interaction with a tape target. *High Energy Density Physics*, 37: 100847, 2020. DOI: `10.1016/j.hedp.2020.100847` (see p. 40)

[93]   L. Obst *et al.* Efficient laser-driven proton acceleration from cylindrical and planar cryogenic hydrogen jets. *Scientific Reports 2017 7:1*, 7: 1–9, 2017. DOI: `10.1038/s41598-017-10589-3` (see p. 40)

[94]   M. Rehwald *et al.* Ultra-short pulse laser acceleration of protons to 80 MeV from cryogenic hydrogen jets tailored to near-critical density. *Nature Communications*, 14: 4009, 2023. DOI: `10.1038/s41467-023-39739-0` (see p. 40)

[95]   R. J. Shalloo *et al.* Automation and control of laser wakefield accelerators using Bayesian optimization. *Nature Communications 2020 11:1*, 11: 1–8, 2020. DOI: `10.1038/s41467-020-20245-6` (see p. 40)

[96]   F. P. Boody, R. Höpfl, H. Hora, and J. C. Kelly. Laser-driven ion source for reduced-cost implantation of metal ions for strong reduction of dry friction and increased durability. *Laser and Particle Beams*, 14: 443–448, 1996. DOI: `10.1017/S0263034600010132` (see p. 40)

[97]   R. J. Clarke *et al.* Detection of short lived radioisotopes as a fast diagnostic for intense laser-solid interactions. *Applied Physics Letters*, 89: 141117, 2006. DOI: `10.1063/1.2358940` (see pp. 40, 124)

[98]   K. W. D. Ledingham, P. McKenna, and R. P. Singhal. Applications for Nuclear Phenomena Generated by Ultra-Intense Lasers. *Science*, 300: 1107–1111, 2003. DOI: `10.1126/science.1080552` (see p. 40)

[99] Z. Sun. Review: Production of nuclear medicine radioisotopes with ultra-intense lasers. *AIP Advances*, 11: 040701, 2021. DOI: 10.1063/5.0042796 (see p. 41)

[100] M. I. K. Santala *et al.* Production of radioactive nuclides by energetic protons generated from intense laser-plasma interactions. *Applied Physics Letters*, 78: 19–21, 2001. DOI: 10.1063/1.1335849 (see p. 41)

[101] A. Yogo *et al.* Measurement of relative biological effectiveness of protons in human cancer cells using a laser-driven quasimonoenergetic proton beamline. *Applied Physics Letters*, 98: 053701, 2011. DOI: 10.1063/1.3551623 (see pp. 41, 42)

[102] D. Doria *et al.* Biological effectiveness on live cells of laser driven protons at dose rates exceeding 109 Gy/s. *AIP Advances*, 2: 011209, 2012. DOI: 10.1063/1.3699063 (see pp. 41, 42)

[103] F. Hanton *et al.* DNA DSB Repair Dynamics following Irradiation with Laser-Driven Protons at Ultra-High Dose Rates. *Scientific Reports*, 9: 4471, 2019. DOI: 10.1038/s41598−019−40339−6 (see p. 41)

[104] K Zeil *et al.* Dose-controlled irradiation of cancer cells with laser-accelerated proton pulses. *Applied Physics B*, 110: 437–444, 2013. DOI: 10.1007/s00340−012−5275−3 (see p. 42)

[105] S. Raschke *et al.* Ultra-short laser-accelerated proton pulses have similar DNA-damaging effectiveness but produce less immediate nitroxidative stress than conventional proton beams. *Scientific Reports*, 6: 32441, 2016. DOI: 10.1038/srep32441 (see p. 42)

[106] L. Manti *et al.* The radiobiology of laser-driven particle beams: focus on sub-lethal responses of normal human cells. *Journal of Instrumentation*, 12: C03084, 2017. DOI: 10.1088/1748−0221/12/03/C03084 (see p. 42)

[107] P. Antici *et al.* Numerical study of a linear accelerator using laser-generated proton beams as a source. *Journal of Applied Physics*, 104: 124901, 2008. DOI: 10.1063/1.3021160 (see p. 42)

[108] A. Wakita *et al.* Characteristics of a laser-produced proton beam improved by a synchronous RF field. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 599: 15–19, 2009. DOI: 10.1016/j.nima.2008.10.025 (see p. 43)

[109] M. Tabak *et al.* Ignition and high gain with ultrapowerful lasers. *Physics of Plasmas*, 1: 1626–1634, 1994. DOI: 10.1063/1.870664 (see p. 43)

[110] B. Hegelich *et al.* Experimental demonstration of particle energy, conversion efficiency and spectral shape required for ion-based fast ignition. *Nuclear Fusion*, 51: 083011, 2011. DOI: 10.1088/0029−5515/51/8/083011 (see p. 43)

[111] J. J. Honrubia, J. C. Fernández, M. Temporal, B. M. Hegelich, and J. Meyer-ter-Vehn. Fast ignition of inertial fusion targets by laser-driven carbon beams. *Physics of Plasmas*, 16: 102701, 2009. DOI: 10.1063/1.3234248 (see p. 43)

[112] J. J. Honrubia and M. Murakami. Ion beam requirements for fast ignition of inertial fusion targets. *Physics of Plasmas*, 22: 012703, 2015. DOI: 10.1063/1.4905904 (see pp. 43, 44)

[113] M. Temporal, J. J. Honrubia, and S. Atzeni. Proton-beam driven fast ignition of inertially confined fuels: Reduction of the ignition energy by the use of two proton beams with radially shaped profiles. *Physics of Plasmas*, 15: 052702, 2008. DOI: 10.1063/1.2918316 (see p. 44)

[114] V. Y. Bychenkov *et al.* Pion production under the action of intense ultrashort laser pulse on a solid target. *Journal of Experimental and Theoretical Physics Letters*, 74: 586–589, 2001. DOI: 10.1134/1.1455066 (see p. 44)

[115] S. Bulanov *et al.* Neutrino oscillation studies with laser-driven beam dump facilities. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 540: 25–41, 2005. DOI: https://doi.org/10.1016/j.nima.2004.11.013 (see p. 44)

[116] M. P. Stockli and T. Nakagawa. Ion Injectors for High-Intensity Accelerators. *Reviews of Accelerator Science and Technology*, 06: 197–219, 2013. DOI: 10.1142/S1793626813300090 (see p. 45)

[117] N. Chauvin *et al.* " Source and Injector Design for Intense Light Ion Beams Including Space Charge Neutralisation" in *25th International Linear Accelerator Conference*. 2011. TH302 (see p. 45)

[118] M. P. Stockli *et al.* Recent performance of the SNS H- ion source and low-energy beam transport system. *Review of Scientific Instruments*, 85: 02B137, 2014. DOI: 10.1063/1.4862205 (see p. 45)

[119] L. Arnaudon *et al.* Linac4 Technical Design Report. techreport revised version submitted on 2006-12-14 09:00:40 Geneva: CERN, 2006 (see p. 45)

[120] U. Akira. J-PARC H- ion source and space-charge neutralized LEBT for 100-mA high energy and high duty factor LINACs. *Review of Scientific Instruments*, 91: 033312, 2020. DOI: 10.1063/1.5126365 (see p. 45)

[121] H. Li *et al.* "The Status of CSNS Front End" in *Proc. 29th Linear Accelerator Conference (LINAC'18), Beijing, China, 16-21 September 2018*: (Beijing, China). Linear Accelerator Conference 29 Geneva, Switzerland: JACoW Publishing, january 2019. 771–773 DOI: doi:10.18429/JACoW-LINAC2018-THPO039 (see p. 45)

[122] O. Willi *et al.* Laser triggered micro-lens for focusing and energy selection of MeV protons. *Laser and Particle Beams*, 25: 71–77, 2007. DOI: 10.1017/S0263034607070115 (see p. 46)

[123] S. Kar *et al.* Guided post-acceleration of laser-driven ions by a miniature modular structure. *Nature Communications 2016 7:1*, 7: 1–7, 2016. DOI: 10.1038/ncomms10792 (see p. 46)

[124] H. Ahmed *et al.* High energy implementation of coil-target scheme for guided re-acceleration of laser-driven protons. *Scientific Reports 2021 11:1*, 11: 1–7, 2021. DOI: 10.1038/s41598-020-77997-w (see p. 46)

[125] A. W. Chao, K. H. Mess, M. Tigner, and F. Zimmermann. *Handbook of Accelerator Physics and Engineering*. 2nd ed. WORLD SCIENTIFIC, 2013. DOI: 10.1142/8543 (see pp. 46, 61, 153)

[126] L. R. Prost. Selected List of Low Energy Beam Transport Facilities for Light-Ion, High-Intensity Accelerators. 2016. DOI: 10.48550/ARXIV.1602.05488 (see p. 47)

[127] M. Reiser. *Review of Charged Particle Dynamics*. John Wiley & Sons, Ltd, 2008. DOI: 10.1002/9783527622047.ch2 (see pp. 47, 52, 156, 216, 219)

[128] I. Hofmann. Performance of solenoids versus quadrupoles in focusing and energy selection of laser accelerated protons. *Phys. Rev. ST Accel. Beams*, 16: 041302, 2013. DOI: 10.1103/PhysRevSTAB.16.041302 (see p. 48)

[129] W. K. H. Panofsky and W. R. Baker. A Focusing Device for the External 350-Mev Proton Beam of the 184-Inch Cyclotron at Berkeley. *Review of Scientific Instruments*, 21: 445–447, 1950. DOI: 10.1063/1.1745611 (see p. 50)

[130] J. van Tilborg *et al.* Active Plasma Lensing for Relativistic Laser-Plasma-Accelerated Electron Beams. *Phys. Rev. Lett.*, 115: 184802, 2015. DOI: 10.1103/PhysRevLett.115.184802 (see p. 51)

[131] R. Pompili *et al.* Experimental characterization of active plasma lensing for electron beams. *Applied Physics Letters*, 110: 104101, 2017. DOI: 10.1063/1.4977894 (see p. 51)

[132] C. A. Lindstrøm *et al.* Emittance Preservation in an Aberration-Free Active Plasma Lens. *Phys. Rev. Lett.*, 121: 194801, 2018. DOI: 10.1103/PhysRevLett.121.194801 (see p. 51)

[133] J. van Tilborg *et al.* Nonuniform discharge currents in active plasma lenses. *Phys. Rev. Accel. Beams*, 20: 032803, 2017. DOI: 10.1103/PhysRevAccelBeams.20.032803 (see p. 51)

[134] C. A. Lindstrøm and E. Adli Analytic plasma wakefield limits for active plasma lenses 2018 DOI: 10.48550/ARXIV.1802.02750 (see p. 51)

[135] R. Pompili *et al.* Focusing of High-Brightness Electron Beams with Active-Plasma Lenses. *Phys. Rev. Lett.*, 121: 174801, 2018. DOI: 10.1103/PhysRevLett.121.174801 (see p. 51)

[136] K. N. Sjobak *et al.* Strong focusing gradient in a linear active plasma lens. *Phys. Rev. Accel. Beams*, 24: 121306, 2021. DOI: 10.1103/PhysRevAccelBeams.24.121306 (see p. 51)

[137] T. Yang *et al.* Designing of active plasma lens for focusing laser-plasma-accelerated pulsed proton beams. *Phys. Rev. Accel. Beams*, 24: 031301, 2021. DOI: 10.1103/PhysRevAccelBeams.24.031301 (see p. 51)

[138] D. S. Chang, F. D. Lasley, I. J. Das, M. S. Mendonca, J. R. Dynlacht, *et al. Basic radiotherapy physics and biology*. Springer, 2014. (see p. 56)

[139] W. D. Bloomer and S. Hellman. Normal Tissue Responses to Radiation Therapy. *New England Journal of Medicine*, 293: PMID: 1093032, 80–83, 1975. DOI: `10.1056/NEJM197507102930206` (see p. 55)

[140] C. Chargari *et al.* Optimize and refine therapeutic index in radiation therapy: Overview of a century. *Cancer Treatment Reviews*, 45: 58–67, 2016. DOI: `10.1016/j.ctrv.2016.03.001` (see p. 56)

[141] L. J. Nevay *et al.* BDSIM: An accelerator tracking code with particle-matter interactions. *Computer Physics Communications*, 107200, 2020. (see pp. 57, 101, 108, 180, 243)

[142] H. Paganetti and P. van Luijk. Biological Considerations When Comparing Proton Therapy With Photon Therapy. *Seminars in Radiation Oncology*, 23: Controversies in Proton Therapy, 77–87, 2013. DOI: `10.1016/j.semradonc.2012.11.002` (see p. 58)

[143] H. Paganetti. Relative biological effectiveness (RBE) values for proton beam therapy. Variations as a function of biological endpoint, dose, and linear energy transfer. *Physics in Medicine & Biology*, 59: R419, 2014. DOI: `10.1088/0031-9155/59/22/R419` (see p. 58)

[144] B. Jones, S. McMahon, and K. Prise. The Radiobiology of Proton Therapy: Challenges and Opportunities Around Relative Biological Effectiveness. *Clinical Oncology*, 30: Proton Beam and Particle Therapy, 285–292, 2018. DOI: `10.1016/j.clon.2018.01.010` (see p. 58)

[145] H. Paganetti *et al.* Report of the AAPM TG-256 on the relative biological effectiveness of proton beams in radiation therapy. *Medical Physics*, 46: e53–e78, 2019. DOI: `10.1002/mp.13390` (see p. 58)

[146] C. P. Karger and P. Peschke. RBE and related modeling in carbon-ion therapy. *Physics in Medicine & Biology*, 63: 01TR02, 2017. DOI: `10.1088/1361-6560/aa9102` (see p. 58)

[147] R. R. Wilson. Radiological Use of Fast Protons. *Radiology*, 47: PMID: 20274616, 487–491, 1946. DOI: `10.1148/47.5.487` (see p. 58)

[148] M. Durante and J. Flanz. Charged particle beams to cure cancer: Strengths and challenges. *Seminars in Oncology*, 46: 219–225, 2019. DOI: `10.1053/j.seminoncol.2019.07.007` (see p. 59)

[149] S. Girdhani, R. Sachs, and L. Hlatky. Biological Effects of Proton Radiation: What We Know and Don't Know. *Radiation Research*, 179: 257 –272, 2013. DOI: `10.1667/RR2839.1` (see p. 59)

[150] D. B. Eekers *et al.* Benefit of particle therapy in re-irradiation of head and neck patients. Results of a multicentric in silico ROCOCO trial. *Radiotherapy and Oncology*, 121: 387–394, 2016. DOI: `10.1016/j.radonc.2016.08.020` (see p. 59)

[151] H. X. Q. Norman *et al.* Review of Technologies for Ion Therapy Accelerators 2021. DOI: `10.48550/ARXIV.2105.08103` (see p. 59)

[152] J. R. Castro. Results of heavy ion radiotherapy. *Radiation and Environmental Biophysics*, 34: 45–48, 1995. DOI: `10.1007/BF01210545/METRICS` (see p. 59)

[153] M. Krämer *et al.* Helium ions for radiotherapy? Physical and biological verifications of a novel treatment modality. *Medical Physics*, 43: 1995–2004, 2016. DOI: `10.1118/1.4944593` (see p. 59)

[154] B. Knäusl, H. Fuchs, K. Dieckmann, and D. Georg. Can particle beam therapy be improved using helium ions? – a planning study focusing on pediatric patients. *Acta Oncologica*, 55: PMID: 26750803, 751–759, 2016. DOI: `10.3109/0284186X.2015.1125016` (see p. 59)

[155] O. Sokol *et al.* Oxygen beams for therapy: advanced biological treatment planning and experimental verification. *Physics in Medicine & Biology*, 62: 7798, 2017. DOI: `10.1088/1361-6560/aa88a0` (see p. 59)

[156] T. Tessonnier *et al.* Proton and helium ion radiotherapy for meningioma tumors: a Monte Carlo-based treatment planning comparison. *Radiation Oncology*, 13: 2, 2018. DOI: `10.1186/s13014-017-0944-3` (see p. 59)

[157] M. Vretenar *et al.* "A Compact Synchrotron for Advanced Cancer Therapy with Helium and Proton Beams" in *Proc. IPAC'22*: (Bangkok, Thailand). International Particle Accelerator Conference 13 JACoW Publishing, Geneva, Switzerland, july 2022. TUOZGD2, 811–814 DOI: `10.18429/JACoW-IPAC2022-TUOZGD2` (see p. 59)

[158] R. Mohan and D. Grosshans. Proton therapy – Present and future. *Advanced Drug Delivery Reviews*, 109: 26–44, 2017. DOI: `10.1016/j.addr.2016.11.006` (see p. 59)

[159] E. O. Lawrence and M. S. Livingston. The Production of High Speed Light Ions Without the Use of High Voltages. *Phys. Rev.*, 40: 19–35, 1932. DOI: `10.1103/PhysRev.40.19` (see p. 60)

[160] Mevion Medical Systems `https://www.mevion.com/products/mevion-s250-proton-therapy-system`. (see p. 60)

[161] H. Owen, A. Lomax, and S. Jolly. Current and future accelerator technologies for charged particle therapy. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 809: 96–104, 2016. DOI: `10.1016/j.nima.2015.08.038` (see p. 60)

[162] V. Smirnov and S. Vorozhtsov. Modern compact accelerators of cyclotron type for medical applications. *Physics of Particles and Nuclei*, 47: 863–883, 2016. DOI: `10.1134/S1063779616050051` (see p. 60)

[163] Y. Jongen *et al.* Compact superconducting cyclotron C400 for hadron therapy. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 624: 47 –53, 2010. DOI: `10.1016/j.nima.2010.09.028` (see p. 60)

[164] J. Flanz. Accelerators for charged particle therapy. *Modern Physics Letters A*, 30: 1540020, 2015. DOI: `10.1142/S0217732315400209` (see p. 61)

[165] A. Degiovanni and U. Amaldi. History of hadron therapy accelerators. *Physica Medica*, 31: 322–332, 2015. DOI: `10.1016/j.ejmp.2015.03.002` (see p. 61)

[166] J. M. Schippers, A. Lomax, A. Garonna, and K. Parodi. Can Technological Improvements Reduce the Cost of Proton Radiation Therapy? *Seminars in Radiation Oncology*, 28: 150–159, 2018. DOI: `10.1016/j.semradonc.2017.11.007` (see p. 61)

[167] P. Montay-Gruel *et al.* Irradiation in a flash: Unique sparing of memory in mice after whole brain irradiation with dose rates above 100Gy/s. *Radiotherapy and Oncology*, 124: 15th International Wolfsberg Meeting 2017, 365–369, 2017. DOI: `.1016/j.radonc.2017.05.003` (see p. 62)

[168] B. W. Loo *et al.* (P003) Delivery of Ultra-Rapid Flash Radiation Therapy and Demonstration of Normal Tissue Sparing After Abdominal Irradiation of Mice. *International Journal of Radiation Oncology, Biology, Physics*, 98: Proceedings of the American Radium Society, E16, 2017. DOI: `10.1016/j.ijrobp.2017.02.101` (see p. 62)

[169] J. Bourhis *et al.* Treatment of a first patient with FLASH-radiotherapy. *Radiotherapy and Oncology*, 139: FLASH radiotherapy International Workshop, 18–22, 2019. DOI: `10.1016/j.radonc.2019.06.019` (see p. 62)

[170] A. E. Mascia *et al.* Proton FLASH Radiotherapy for the Treatment of Symptomatic Bone Metastases: The FAST-01 Nonrandomized Trial. *JAMA Oncology*, 2022. DOI: `10.1001/jamaoncol.2022.5843` (see p. 62)

[171] M.-C. Vozenin, J. Hendry, and C. Limoli. Biological Benefits of Ultra-high Dose Rate FLASH Radiotherapy: Sleeping Beauty Awoken. *Clinical Oncology*, 31: 407–415, 2019. DOI: `10.1016/j.clon.2019.04.001` (see p. 62)

[172] C. Billena and A. J. Khan. A Current Review of Spatial Fractionation: Back to the Future? *International Journal of Radiation Oncology, Biology, Physics*, 104: 177–187, 2019. DOI: `10.1016/j.ijrobp.2019.01.073` (see p. 63)

[173] F. A. Dilmanian *et al.* Response of rat intracranial 9L gliosarcoma to microbeam radiation therapy. *Neuro-Oncology*, 4: 26–38, 2002. DOI: `10.1093/NEUONC/4.1.26` (see p. 63)

[174] A. Bouchet *et al.* Better Efficacy of Synchrotron Spatially Microfractionated Radiation Therapy Than Uniform Radiation Therapy on Glioma. *International Journal of Radiation Oncology, Biology, Physics*, 95: 1485–1494, 2016. DOI: `10.1016/j.ijrobp.2016.03.040` (see p. 63)

[175] M. Sammer *et al.* Proton pencil minibeam irradiation of an in-vivo mouse ear model spares healthy tissue dependent on beam size. *PLOS ONE*, 14: 1–15, 2019. DOI: `10.1371/journal.pone.0224873` (see p. 63)

[176] I. Martínez-Rovira, W. González, S. Brons, and Y. Prezado. Carbon and oxygen minibeam radiation therapy: An experimental dosimetric evaluation. *Medical Physics*, 44: 4223–4229, 2017. DOI: `10.1002/mp.12383` (see p. 63)

[177] T. Schneider. *Advancing the generation of proton minibeams for radiation therapy*. Theses. Université Paris-Saclay, december 2020. (see pp. 63, 64, 67)

[178] S. Kanagavelu *et al.* In Vivo Effects of Lattice Radiation Therapy on Local and Distant Lung Cancer: Potential Role of Immunomodulation. *Radiation Research*, 182: 149 –162, 2014. DOI: `10.1667/RR3819.1` (see p. 63)

[179] A. Bouchet, R. Serduc, J. A. Laissue, and V. Djonov. Effects of microbeam radiation therapy on normal and tumoral blood vessels. *Physica Medica*, 31: Radiation Therapy with Synchrotron Radiation: Achievements and Challenges, 634–641, 2015. DOI: `10.1016/j.ejmp.2015.04.014` (see p. 63)

[180] A. Peeters *et al.* How costly is particle therapy? Cost analysis of external beam radiotherapy with carbon-ions, protons and photons. *Radiotherapy and Oncology*, 95: 45–53, 2010. DOI: `10.1016/j.radonc.2009.12.002` (see p. 64)

[181] G. Aymar *et al.* The Laser-hybrid Accelerator for Radiobiological Applications 2020. DOI: `10.48550/ARXIV.2006.00493` (see p. 65)

[182] M. Noaman-ul-Haq *et al.* Statistical analysis of laser driven protons using a high-repetition-rate tape drive target system. *Phys. Rev. Accel. Beams*, 20: 041301, 2017. DOI: `10.1103/PhysRevAccelBeams.20.041301` (see p. 66)

[183] K. R. Symon, D. W. Kerst, L. W. Jones, L. J. Laslett, and K. M. Terwilliger. Fixed-Field Alternating-Gradient Particle Accelerators. *Phys. Rev.*, 103: 1837–1859, 1956. DOI: `10.1103/PhysRev.103.1837` (see p. 66)

[184] M. Tanigaki *et al.* "Present status of FFAG accelerators in KURRI for ADS study" in *Proceedings of EPAC 2006*. Citeseer 2006. 2367 (see p. 67)

[185] K. Peach *et al.* "PAMELA - a model for an FFAG based hadron therapy machine" in *2007 IEEE Particle Accelerator Conference (PAC)*. 2007. 2880–2882 DOI: `10.1109/PAC.2007.4440607` (see p. 67)

[186] J. Pozimski and M. Aslaninejad. Gabor lenses for capture and energy selection of laser driven ion beams in cancer treatment. *Laser and Particle Beams*, 31: 723–733, 2013. DOI: `10.1017/S0263034613000761` (see pp. 69, 91)

[187] J. R. Danielson, D. H. E. Dubin, R. G. Greaves, and C. M. Surko. Plasma and trap-based techniques for science with positrons. *Rev. Mod. Phys.*, 87: 247–306, 2015. DOI: `10.1103/RevModPhys.87.247` (see pp. 71–76, 78, 90, 100, 209, 210, 262)

[188] J. H. Malmberg and C. F. Driscoll. Long-Time Containment of a Pure Electron Plasma. *Phys. Rev. Lett.*, 44: 654–657, 1980. DOI: `10.1103/PhysRevLett.44.654` (see pp. 71, 79, 123)

[189] D. H. E. Dubin and T. M. O'Neil. Trapped nonneutral plasmas, liquids, and crystals (the thermal equilibrium states). *Rev. Mod. Phys.*, 71: 87–172, 1999. DOI: 10.1103/RevModPhys.71.87 (see pp. 71, 77)

[190] R. C. Davidson. *Physics of nonneutral plasmas*. Imperial College Press London, 2001. (see pp. 75–78, 96, 235, 237, 259)

[191] A. J. Theiss, R. A. Mahaffey, and A. W. Trivelpiece. Rigid-Rotor Equilibria of Nonneutral Plasmas. *Phys. Rev. Lett.*, 35: 1436–1438, 1975. DOI: 10.1103/PhysRevLett.35.1436 (see p. 75)

[192] L. Brillouin. A Theorem of Larmor and Its Importance for Electrons in Magnetic Fields. *Phys. Rev.*, 67: 260–266, 1945. DOI: 10.1103/PhysRev.67.260 (see pp. 75, 83)

[193] C. F. Driscoll, J. H. Malmberg, and K. S. Fine. Observation of transport to thermal equilibrium in pure electron plasmas. *Phys. Rev. Lett.*, 60: 1290–1293, 1988. DOI: 10.1103/PhysRevLett.60.1290 (see p. 77)

[194] S. A. Prasad and T. M. O'Neil. Finite length thermal equilibria of a pure electron plasma column. *The Physics of Fluids*, 22: 278–281, 1979. DOI: 10.1063/1.862578 (see pp. 77, 79, 259)

[195] J. Notte and J. Fajans. The effect of asymmetries on non-neutral plasma confinement time. *Physics of Plasmas*, 1: 1123–1127, 1994. DOI: 10.1063/1.870762 (see pp. 79, 89, 90)

[196] K. S. Fine. *Experiments with the l=1 Diocotron Mode*. Ph. D. Thesis. University of California San Diego, 1988. (see p. 79)

[197] C. F. Driscoll and J. H. Malmberg. Length-Dependent Containment of a Pure Electron-Plasma Column. *Phys. Rev. Lett.*, 50: 167–170, 1983. DOI: 10.1103/PhysRevLett.50.167 (see pp. 79, 89, 90, 134)

[198] J. M. Kriesel and C. F. Driscoll. Two Regimes of Asymmetry-Induced Transport in Non-neutral Plasmas. *Phys. Rev. Lett.*, 85: 2510–2513, 2000. DOI: 10.1103/PhysRevLett.85.2510 (see p. 79)

[199] PulsarPhysics. General Particle Tracer. http://www.pulsar.nl/gpt/index.html. (see pp. 81, 180, 190)

[200] K Schulte *et al.* "Studies on electron cloud dynamics for an optimized space charge lens design" in *Proc. IPAC'11*. JACoW Publishing, Geneva, Switzerland, 2011. THPS034, 3493–3495 (see pp. 83, 87)

[201] M. Droba *et al.* "High-current beam transport simulations including Gabor lenses in varying on-neutral plasma states" in *Proceedings of the 4th International Particle Accelerator Conference (IPAC 2013)*: (Shanghai, China). JACoW Publishing, Geneva, Switzerland, 2013. TUPWO008, 1892–1894 (see pp. 83, 87)

[202] R. H. Levy. Diocotron Instability in a Cylindrical Geometry. *The Physics of Fluids*, 8: 1288–1295, 1965. DOI: 10.1063/1.1761400 (see p. 83)

[203] W. Knauer. Diocotron Instability in Plasmas and Gas Discharges. *Journal of Applied Physics*, 37: 602–611, 1966. DOI: 10.1063/1.1708223 (see p. 83)

[204] R. L. Kyhl and H. F. Webster. Breakup of hollow cylindrical electron beams. *IRE Transactions on Electron Devices*, 3: 172–183, 1956. (see p. 83)

[205] K. Schulte *et al.* Electron cloud dynamics in a gabor space charge lens. *IPAC 2012 - International Particle Accelerator Conference 2012*, 0: 1164–1166, 2012. (see p. 83)

[206] H. Lefevre, R. Connolly, D. Xi, G. Sieger, and J. Overley. Can an electron plasma lens produce sub-micrometer size focal spots of MeV ions? *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 10-11: 707–712, 1985. DOI: 10.1016/0168-583X(85)90089-8 (see pp. 84, 85, 89, 210)

[207] G. E. Sieger, R. C. Connolly, and H. W. Lefevre. Dynamic concentration of positive ions near the axis of an electron plasma column. *The Physics of Fluids*, 27: 291–295, 1984. DOI: 10.1063/1.864481 (see p. 84)

[208] J. A. Palkovic, F. E. Mills, C. Schmidt, and D. E. Young. Gabor lens focusing of a proton beam. *Review of Scientific Instruments*, 61: 550–552, 1990. DOI: 10.1063/1.1141248 (see p. 85)

[209] Q. Hong, Y. Qingchang, and H. Jiachang. Characteristics of self-sustaining magnetically confined electron clouds under different electromagnetic fields and pressures. *Physica Scripta*, 48: 474, 1993. DOI: 10.1088/0031-8949/48/4/016 (see pp. 85, 210)

[210] K. Thoma, M. Droba, and O. Meusel. "Investigation, Simulation and First Measurements of a 2 m Long Electron Column Trapped in a Gabor Lens Device" in *Proc. IPAC'22*: (Bangkok, Thailand). International Particle Accelerator Conference 13 JACoW Publishing, Geneva, Switzerland, july 2022. WEPOTK002, 2023–2026 DOI: 10.18429/JACoW-IPAC2022-WEPOTK002 (see p. 85)

[211] O. Meusel. *Fokussierung und Transport von Ionenstrahlen mit Raumladungslinsen*. PhD Thesis. Johann Wolfgang Goethe-Universität in Frankfurt am Main, 2006. (see p. 86)

[212] E. M. Hollmann, F. Anderegg, and C. F. Driscoll. Confinement and manipulation of non-neutral plasmas using rotating wall electric fields. *Physics of Plasmas*, 7: 2776–2789, 2000. DOI: 10.1063/1.874128 (see pp. 89, 129, 205)

[213] G. W. Hart and B. G. Peterson. Finding the radial parallel temperature profile in a non-neutral plasma using equilibrium calculations on experimental data. *Physics of Plasmas*, 13: 022101, 2006. DOI: 10.1063/1.2167586 (see p. 89)

[214] J. R. Danielson, T. R. Weber, and C. M. Surko. Plasma manipulation techniques for positron storage in a multicell trap. *Physics of Plasmas*, 13: 123502, 2006. DOI: 10.1063/1.2390690 (see p. 89)

[215] J. R. Danielson and C. M. Surko. Radial compression and torque-balanced steady states of single-component plasmas in Penning-Malmberg traps. *Physics of Plasmas*, 13: 055706, 2006. DOI: `10.1063/1.2179410` (see pp. 89, 205)

[216] J. Fajans and C. M. Surko. Plasma and trap-based techniques for science with antimatter. *Physics of Plasmas*, 27: 030601, 2020. DOI: `10.1063/1.5131273` (see pp. 90, 205)

[217] K. S. Fine, C. F. Driscoll, and J. H. Malmberg. Measurements of a nonlinear diocotron mode in pure electron plasmas. *Phys. Rev. Lett.*, 63: 2232–2235, 1989. DOI: `10.1103/PhysRevLett.63.2232` (see pp. 90, 100, 113)

[218] L. Turner. Brillouin limit for non-neutral plasma in inhomogeneous magnetic fields. *Physics of Fluids B: Plasma Physics*, 3: 1355–1363, 1991. DOI: `10.1063/1.859700` (see p. 90)

[219] E. H. Chao, R. C. Davidson, S. F. Paul, and K. A. Morrison. Effects of background gas pressure on the dynamics of a nonneutral electron plasma confined in a Malmberg-Penning trap. *Physics of Plasmas*, 7: 831–838, 2000. DOI: `10.1063/1.873879` (see p. 90)

[220] J. Pozimski, M Aslaninejad, P. Posocco, *et al.* "Advanced Gabor lens lattice for medical applications" in *Proceedings of IPAC, Shanghai, China*. Citeseer, 2013. 1277–1279 (see p. 91)

[221] J. Pozimski *et al.* "A Ready-to-use Application of Laser-Plasma Accelerators using Gabor Lenses" in *Proceedings of The North American Particle Accelerator Conference, NA-PAC2013, Pasadena, CA USA*. `http://accelconf.web.cern.ch/AccelConf/PAC2013/papers/thpsm12.pdf` 2013. (see p. 91)

[222] P. Posocco, M. Merchant, J. Pozimski, and Y. Xia. "First Test of The Imperial College Gabor (Plasma) Lens prototype at the Surrey Ion Beam centre" in *Proceedings of International Particle Accelerator Conference (IPAC'16), Busan, Korea, May 8-13, 2016*. 7 Geneva, Switzerland: JACoW, june 2016. 1598–1600 DOI: `10.18429/JACoW-IPAC2016-TUPMY024` (see pp. 92, 99, 201)

[223] D. C. Meeker Finite Element Method Magnetics, Version 4.2 (21Apr2019 Build). `https://www.femm.info` (see pp. 93, 95, 171–173, 243, 250)

[224] University of Surrey Ion Beam Centre `https://www.surrey.ac.uk/ion-beam-centre`. Accessed: 2023-04-10 (see pp. 94, 98)

[225] G. Lapenta. Kinetic plasma simulation: Particle in cell method. *XII Carolus Magnus Summer School on Plasma and Fusion Energy Physics*, 76–85, 2015. (see pp. 94, 256, 257)

[226] M. Reiser. Comparison of Gabor Lens, Gas Focusing, and Electrostatic Quadrople Focusing for Low Energy Ion Beams. *Conf. Proc. C*, 8903201: 1744, 1989. (see pp. 95, 209)

[227] Electron-Impact Cross Sections for Ionization and Excitation, NIST `https://physics.nist.gov/PhysRefData/Ionization/molTable.html` 2020 (see p. 95)

[228] Y. Itikawa and N. Mason. Cross Sections for Electron Collisions with Water Molecules. *Journal of Physical and Chemical Reference Data*, 34: 1–22, 2005. DOI: 10.1063/1.1799251 (see p. 95)

[229] Scientific Computing Resources for Facilities https://www.scarf.rl.ac.uk 2023 (see p. 95)

[230] K. Schulte, M. Droba, O. Meusel, and U. Ratzinger. "Investigation of diagnostic techniques on a nonneutral plasma" in *Proceedings of the 10th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators DIPAC2011, Hamburg, Germany, 2011*. MOPD93 2011. 266 –268 (see p. 96)

[231] R. C. Davidson and G. M. Felice. Influence of profile shape on the diocotron instability in a non-neutral plasma column. *Physics of Plasmas*, 5: 3497–3511, 1998. DOI: 10.1063/1.873067 (see pp. 100, 235, 236)

[232] R. W. Gould. Dynamics of non-neutral plasmas. *Physics of Plasmas*, 2: 2151–2163, 1995. DOI: 10.1063/1.871302 (see pp. 100, 235, 236)

[233] Tech-X Corporation. VSim: an electromagnetics and plasma computational application. https://www.txcorp.com/vsim. Accessed: 2023-02-15 (see pp. 100, 145, 251)

[234] A. J. Peurrung and J. Fajans. Experimental dynamics of an annulus of vorticity in a pure electron plasma. *Physics of Fluids A: Fluid Dynamics*, 5: 493–499, 1993. DOI: 10.1063/1.858872 (see p. 102)

[235] K. S. Fine and C. F. Driscoll. The finite length diocotron mode. *Physics of Plasmas*, 5: 601–607, 1998. DOI: 10.1063/1.872752 (see pp. 110, 138)

[236] N. C. Hurst, J. R. Danielson, C. J. Baker, and C. M. Surko. Finite-length, large-amplitude diocotron mode dynamics. *AIP Conference Proceedings*, 1668: 020003, 2015. DOI: 10.1063/1.4923106 (see pp. 113, 123, 129)

[237] M. Knoop, N. Madsen, and R. C. Thompson. *Physics with Trapped Charged Particles*. Imperial College Press, 2014. DOI: 10.1142/p928 (see p. 121)

[238] C. Evans. *Design and implementation of a third stage positron accumulator device*. MSc Thesis. Swansea University, december 2014. (see p. 124)

[239] A. Deller. *Positron accumulation and laser excitation of the positronium atom*. phdthesis. Swansea Univeristy, 2013. (see pp. 124, 128, 132)

[240] J. R. Danielson and C. M. Surko. Radial compression and torque-balanced steady states of single-component plasmas in Penning-Malmberg traps. *Physics of Plasmas*, 13: 055706, 2006. DOI: 10.1063/1.2179410 (see pp. 129, 134, 205)

[241] T. Savage. *Towards non-destructive diagnostics oftrapped positrons*. MSc Thesis. Swansea University, january 2019. (see pp. 131, 148)

[242]   C. A. Isaac. *Axialisation of particles in a Penning-type trap by the application of a rotating dipole electric field and its application to positron accumulation.* phdthesis. Swansea Univeristy, 2010. (see p. 132)

[243]   C. Nieter and J. R. Cary. VORPAL: a versatile plasma simulation code. *Journal of Computational Physics*, 196: 448–473, 2004. DOI: 10.1016/j.jcp.2003.11.004 (see p. 145)

[244]   Tech-X Corporation. VSim Reference Manual. https://txcorp.com/images/docs/vsim/latest/VSimReferenceManual/VSimReferenceManual.html. Accessed: 2023-04-10 (see pp. 91, 146, 253, 255)

[245]   D. Tomamsini. Practical Definitions & Formulae for Normal Conducting Magnets. CERN Internal Note 2011-18, EDMS Nr: 1162401. 2011 (see pp. 154, 161, 163, 174)

[246]   T. Zickler. Basic design and engineering of normal-conducting, iron-dominated electromagnets. 2011. DOI: 10.48550/ARXIV.1103.1119 (see pp. 154, 161, 163, 164, 174)

[247]   H. T. Lau. *Medical Applications for Particle Physics*. Imperial College London. 2022. (see pp. 154, 155, 187)

[248]   T. E. Cowan *et al.* Ultralow Emittance, Multi-MeV Proton Beams from a Laser Virtual-Cathode Plasma Accelerator. *Phys. Rev. Lett.*, 92: 204801, 2004. DOI: 10.1103/PhysRevLett.92.204801 (see pp. 155, 156)

[249]   S. S. Nagaitsev and A. Shemyakin. Beam Emittance Calculation in the Presence of an Axially Symmetric Magnetic Field. FERMILAB-TM-2107 Batavia, IL, 2000. DOI: 10.2172/754860 (see p. 156)

[250]   V. Toivanen, T. Kalvas, H. Koivisto, J. Komppula, and O. Tarvainen. Double einzel lens extraction for the JYFL 14 GHz ECR ion source designed with IBSimu. *Journal of Instrumentation*, 8: P05003–P05003, 2013. DOI: 10.1088/1748-0221/8/05/p05003 (see p. 156)

[251]   K. Harres *et al.* Beam collimation and transport of quasineutral laser-accelerated protons by a solenoid field. *Physics of Plasmas*, 17: 023107, 2010. DOI: 10.1063/1.3299391 (see pp. 156, 174)

[252]   C. Yao and Y. Ma. Superconducting materials: Challenges and opportunities for large-scale applications. *iScience*, 24: 2021. DOI: 10.1016/J.ISCI.2021.102541 (see pp. 157, 158)

[253]   M. N. Wilson. Superconducting materials for magnets. 1996. DOI: 10.5170/CERN-1996-003.47 (see p. 157)

[254]   R Flükiger. Superconductivity for Magnets. Comments: 21 pages, contribution to the CAS-CERN Accelerator School: Superconductivity for Accelerators, Erice, Italy, 24 April - 4 May 2013, edited by R. Bailey, 21 p, 2014. DOI: 10.5170/CERN-2014-005.247 (see p. 158)

[255]   A. Yamamoto. Advances in MgB_2 Superconductor Applications for Particle Accelerators. 2022. DOI: 10.48550/ARXIV.2201.09501 (see p. 158)

[256] Applied Superconductivity Center at NHMFL. Engineering critical current density vs. applied field for superconductors available in long lengths. `https://nationalmaglab.org/magnet-development/applied-superconductivity-center/plots`. Online; accessed 11 August 2022 (see p. 158)

[257] H. Kanithi *et al.* Production Results of 11.75 Tesla Iseult/INUMAC MRI Conductor at Luvata. *IEEE Transactions on Applied Superconductivity*, 24: 1–4, 2014. DOI: `10.1109/TASC.2013.2281417` (see p. 158)

[258] M. S. A. Hossain, C Senatore, M Rindfleisch, and R Flükiger. Improvement of Jc by cold high pressure densification of binary, 18-filament in situ MgB_2 wires. *Superconductor Science and Technology*, 24: 075013, 2011. DOI: `10.1088/0953-2048/24/7/075013` (see p. 158)

[259] G. Z. Li *et al.* The critical current density of advanced internal-Mg-diffusion-processed MgB_2 wires. *Superconductor Science and Technology*, 25: 115023, 2012. DOI: `10.1088/0953-2048/25/11/115023` (see p. 158)

[260] T. Boutboul, S. Le Naour, D. Leroy, L. Oberli, and V. Previtali. Critical Current Density in Superconducting Nb-Ti Strands in the 100 mT to 11 T Applied Field Range. *IEEE Transactions on Applied Superconductivity*, 16: 1184–1187, 2006. DOI: `10.1109/TASC.2006.870777` (see p. 158)

[261] A. Beqiri. Design of a superconducting solenoid suitable to work with a laser particle source. 2021. (see p. 158)

[262] J. D. Jackson. *Classical electrodynamics*. 3rd ed. New York, NY: Wiley, 1999. (see p. 159)

[263] G. Brianti and M. Gabriel. Basic expressions for evaluating iron core magnets–a possible procedure to minimize their cost. CERN Note CERN-SI-INT-DL-70-10. 1970 (see pp. 159, 160)

[264] Luvata. Oxygen-free Copper Cu-OF Data Sheet. `https://www.luvata.com/products/hollow-conductors`. Online; accessed 03 August 2022 (see pp. 161, 165, 166)

[265] Engineering ToolBox, (2003). Resistivity and Conductivity - Temperature Coefficients Common Materials. `https://www.engineeringtoolbox.com/resistivity-conductivity-d_418.html`. Online; accessed 11 August 2022 (see p. 161)

[266] Software-Factory H. Schmitz Germany. Pressure Drop Calculator. `http://www.pressure-drop.online/`. Online; accessed 03 August 2022 (see pp. 164, 176–178)

[267] V. Bayliss and J. Boehm. Development of a magnet system for the CERN BDF/SHiP Facility. CERN-SHiP-NOTE-2019-004, 2019. (see p. 165)

[268] J. T. Tanabe. *Iron Dominated Electromagnets*. WORLD SCIENTIFIC, 2005. DOI: `10.1142/5823` (see p. 165)

[269]   G. Sinclair and J. Helms. A review of simple formulae for elastic hoop stresses in cylindrical and spherical pressure vessels: What can be used when. *International Journal of Pressure Vessels and Piping*, 128: 1–7, 2015. DOI: `10.1016/j.ijpvp.2015.01.006` (see p. 165)

[270]   The LhARA Collaboration. Baseline for the LhARA design update. Technical note CCAP-TN-11. `https://ccap.hep.ph.ic.ac.uk/trac/raw-attachment/wiki/Communication/Notes/CCAP-TN-11-LhARA-Design-Baseline.pdf` 2022 (see pp. 174, 180, 189, 192, 193, 197, 198, 200)

[271]   R. L. Spencer, S. N. Rasband, and R. R. Vanfleet. Numerical calculation of axisymmetric non-neutral plasma equilibria. *Physics of Fluids B: Plasma Physics*, 5: 4267–4272, 1993. DOI: `10.1063/1.860594` (see pp. 184, 185)

[272]   R. A. Lane and C. A. Ordonez. Electrostatic equilibria of non-neutral plasmas confined in a Penning trap with axially varying magnetic field. *Physics of Plasmas*, 26: 052511, 2019. DOI: `10.1063/1.5092136` (see p. 184)

[273]   W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, 2007. (see p. 185)

[274]   E. Boella. "The novel Laser-hybrid Accelerator for Radiobiological Applications" Presented at the Laser and Plasma Accelerators Workshop 2023, Lisbon (Portugal), March 6th 2023. `https://lpaw23.tecnico.ulisboa.pt/program` (see pp. 187, 188)

[275]   The LhARA Collaboration. LhARA/ITRF: six month progress report. Report LhARA-Gov-PMB-2023-03. 2023 (see pp. 187, 188)

[276]   S. M. Wiggins *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*. International Society for Optics and Photonics SPIE, 2019. 110360T DOI: `10.1117/12.2520717` (see pp. 187, 188)

[277]   K. I. Thoma, M Droba, and O Meusel. Investigation, simulation and first measurements of a 2 m long electron column trapped in a Gabor-Lens device. *Journal of Physics: Conference Series*, 2420: 012042, 2023. DOI: `10.1088/1742-6596/2420/1/012042` (see p. 204)

[278]   M. Ahmadi *et al.* Enhanced Control and Reproducibility of Non-Neutral Plasmas. *Phys. Rev. Lett.*, 120: 025001, 2018. DOI: `10.1103/PhysRevLett.120.025001` (see p. 205)

[279]   A. Goncharov, A. Dobrovolsky, A. Zatuagan, and I. Protsenko. High-current plasma lens. *IEEE Transactions on Plasma Science*, 21: 573–577, 1993. DOI: `10.1109/27.249645` (see p. 207)

[280]   V. Zhukov, A. Morozov, and T. Shepkin. Experimental investigation plasma focusing of ion beams. *Pismu v ZETF*, 9: `http://jetpletters.ru/ps/1638/article_25020.pdf`, 24, 1969. (see p. 207)

[281] A. Goncharov, A. Zatuagan, and I. Protsenko. Focusing and control of multiaperture ion beams by plasma lenses. *IEEE Transactions on Plasma Science*, 21: 578–581, 1993. DOI: `10.1109/27.249646` (see p. 207)

[282] A. Goncharov *et al.* Some characteristics of moderate energy metal ion beam focusing by a high current plasma lens. *Review of Scientific Instruments*, 69: 1135–1137, 1998. DOI: `10.1063/1.1148647` (see p. 207)

[283] A. A. Goncharov, A. N. Dobrovolsky, I. V. Litovko, I. M. Protsenko, and V. F. Zadorodzny. High-current plasma lens: New results and applications. *Review of Scientific Instruments*, 67: 1155–1157, 1996. DOI: `10.1063/1.1146717` (see p. 207)

[284] A. A. Goncharov, I. M. Protsenko, G. Y. Yushkov, and I. G. Brown. Focusing of high-current, large-area, heavy-ion beams with an electrostatic plasma lens. *Applied Physics Letters*, 75: 911–913, 1999. DOI: `10.1063/1.124551` (see p. 207)

[285] A. Goncharov, I. Protsenko, G. Yushkov, and I. Brown. Manipulating large-area, heavy metal ion beams with a high-current electrostatic plasma lens. *IEEE Transactions on Plasma Science*, 28: 2238–2246, 2000. DOI: `10.1109/27.902252` (see p. 208)

[286] A. Goncharov *et al.* Permanent magnet plasma lens. *Review of Scientific Instruments*, 73: 1001–1003, 2002. DOI: `10.1063/1.1431414` (see p. 208)

[287] A. Goncharov. Status of the electrostatic plasma lens. *Review of Scientific Instruments*, 73: 1004–1006, 2002. DOI: `10.1063/1.1431415` (see p. 208)

[288] A. Goncharov, V. Gorshkov, V. Maslov, V. Zadorozhny, and I. Brown. The low-field permanent magnet electrostatic plasma lens. *Review of Scientific Instruments*, 75: 1662–1664, 2004. DOI: `10.1063/1.1699450` (see p. 208)

[289] Y. Chekh, A. Goncharov, I. Protsenko, and I. G. Brown. Effect of the electrostatic plasma lens on the emittance of a high-current heavy ion beam. *Applied Physics Letters*, 86: 041502, 2005. DOI: `10.1063/1.1855428` (see p. 208)

[290] Y. M. Chekh, A. A. Goncharov, I. M. Protsenko, A. M. Dobrovolsky, and I. G. Brown. Electrostatic plasma lens for accelerator injection application. *Review of Scientific Instruments*, 77: 03B906, 2006. DOI: `10.1063/1.2165577` (see p. 208)

[291] A. Goncharov. Invited Review Article: The electrostatic plasma lens. *Review of Scientific Instruments*, 84: 021101, 2013. DOI: `10.1063/1.4789314` (see p. 208)

[292] J. A. Palkovic, F. E. Mills, C. Schmidt, and D. E. Young. Gabor Lens Focusing of a Negative Ion Beam. *Conf. Proc. C*, 8903201: FERMILAB-CONF-89-115, FNAL-C-89-115, 304, 1989. (see p. 209)

[293] V. P. Goretsky, I. A. Soloshenko, and A. I. Shchedrin. Space-charge lens for focusing negative ion beam. *Review of Scientific Instruments*, 73: 1161–1164, 2002. DOI: `10.1063/1.1431406` (see p. 209)

[294] A. M. Zavalov, V. N. Gorshkov, V. P. Goretskii, and I. A. Soloshenko. Space-charge lens for focusing a negative-ion beam. *Plasma Physics Reports*, 29: 480–484, 2003. DOI: 10.1134/1.1582515 (see p. 209)

[295] I. A. Soloshenko, V. P. Goretsky, V. N. Gorshkov, and A. M. Zavalov. Space charge lens for focusing negative ion beam: Theory and experiment. *Review of Scientific Instruments*, 75: 1774–1776, 2004. DOI: 10.1063/1.1695626 (see p. 209)

[296] D. P. Xi. Density distribution of an electron plasma column in a Penning discharge. *Journal of Physics D: Applied Physics*, 19: 2325, 1986. DOI: 10.1088/0022-3727/19/12/012 (see p. 210)

[297] S. A. Prasad and T. M. O'Neil. Waves in a cold pure electron plasma of finite length. *The Physics of Fluids*, 26: 665–672, 1983. DOI: 10.1063/1.864181 (see p. 210)

[298] A. Wolski. *Beam Dynamics in High Energy Particle Accelerators*. IMPERIAL COLLEGE PRESS, 2014. DOI: 10.1142/p899 (see pp. 211–213, 216)

[299] H. Busch. Berechnung der Bahn von Kathodenstrahlen im axialsymmetrischen elektromagnetischen Felde. *Annalen der Physik*, 386: 974–993, 1926. DOI: 10.1002/andp.19263862507 (see p. 216)

[300] V. Kumar. Understanding the focusing of charged particle beams in a solenoid magnetic field. *American Journal of Physics*, 77: 737–741, 2009. DOI: 10.1119/1.3129242 (see pp. 216, 218)

[301] F. Kroll. *The Study and Development of Pulsed High-field Magnets for Application in Laser-plasma Physics*. Wissenschaftlich-Technischer Bericht / HZDR. 2018. (see p. 218)

[302] C. F. Jekel and G. Venter. *pwlf: A Python Library for Fitting 1D Continuous Piecewise Linear Functions*. https://github.com/cjekel/piecewise_linear_fit_py, 2019. (see p. 229)

[303] F. Anderegg. in: *Chapter 6: Waves in Non-neutral Plasma*. 173–193. 2014. DOI: 10.1142/9781783264063_0006 (see pp. 238, 240)

[304] W. D. White, J. H. Malmberg, and C. F. Driscoll. Resistive-Wall Destabilization of Diocotron Waves. *Phys. Rev. Lett.*, 49: 1822–1826, 1982. DOI: 10.1103/PhysRevLett.49.1822 (see p. 238)

[305] J. M. Dawson. Particle simulation of plasmas. *Rev. Mod. Phys.*, 55: 403–447, 1983. DOI: 10.1103/RevModPhys.55.403 (see p. 251)

[306] R. W. Hockney and J. W. Eastwood. *Computer simulation using particles*. 1988. (see p. 251)

[307] C. K. Birdsall and A. B. Langdon. *Plasma Physics via Computer Simulation*. 1991. (see pp. 251, 257)

[308] Tech-X Corporation. VSim User Guide. https://txcorp.com/images/docs/vsim/latest/VSimUserGuide/VSimUserGuide.html. Accessed: 2023-04-10 (see pp. 252, 253)

[309] M. Knoop, N. Madsen, and R. C. Thompson. *Physics with Trapped Charged Particles*. Imperial College Press, 2014. DOI: 10.1142/p928 (see pp. 254, 256)

[310] H. Fehske, R. Schneider, and A. Weiße. *Computational Many-Particle Physics*. Chapter 6: The Particle-in-Cell Method. Springer Berlin Heidelberg, 2008. DOI: `10.1007/978-3-540-74686-7` (see p. 254)

[311] A. Narimannezhad, C. J. Baker, M. H. Weber, J. Jennings, and K. G. Lynn. Simulation studies of the behavior of positrons in a microtrap with long aspect ratio. *The European Physical Journal D*, 68: 351, 2014. DOI: `10.1140/epjd/e2014-40700-0` (see p. 255)

[312] J. Brackbill. On energy and momentum conservation in particle-in-cell plasma simulation. *Journal of Computational Physics*, 317: 405–427, 2016. DOI: `10.1016/j.jcp.2016.04.050` (see p. 258)