Laser-Ionized Underdense Passive Plasma Lens for Focusing Electron Beams

by

Christopher E. Doss

B.S., West Virginia University, 2016

M.S., University of Colorado Boulder, 2022

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Physics

2023

Committee Members:

Michael Litos, Chair

Iain D. Boyd

Scott Parker

Dmitri Uzdensky

Eric. D. Zimmerman

Doss, Christopher E. (Ph.D., Physics)

Laser-Ionized Underdense Passive Plasma Lens for Focusing Electron Beams

Thesis directed by Prof. Michael Litos

Next generation accelerators and colliders using relativistic electron beams are continuously pushing the demand for higher luminosity beams with smaller and smaller spot sizes. To address this need, passive plasma lenses operating in the nonlinear blowout regime of beam-driven plasma wakefield acceleration (PWFA) are capable of providing focusing forces to electron beams orders of magnitude stronger than conventional quadrupole magnets. To realize these lenses in practice, high intensity lasers can be used to ionize a small volume of gas and producing a plasma lens with precisely determined density profile.

The quality of an electron beam focused by a plasma lens is determined by the phase space evolution of the beam while it is within the plasma wakefield. The strong electric fields can also increase an electron beam's emittance through chromatic phase spreading, which deteriorates the transverse quality of the beam. This dissertation covers the formalism used to describe the focusing of an electron beam from a passive, underdense plasma lens and demonstrates use cases for these lenses in experiments using relativistic electron beams.

To perform plasma lens experiments, we propose a setup to ionize a 100 μm scale plasma lens via laser ionization of a gas jet outflow at the FACET-II accelerator facility of SLAC National Accelerator Laboratory. We investigate possible focusing aberrations induced by nonuniform transverse density profiles. Finally, we report on experimental progress towards the demonstration of an underdense thin plasma lens, including the analysis of preliminary data from commissioning shifts carried out at FACET-II.

Acknowledgements

During the course of my Ph.D. I've had the opportunity to work with amazing people and it would be impossible to fully summarize the entire impact on my life. Foremost, I wish to thank my advisor, Mike Litos, for his continued support and optimism during the long course of my thesis. He demonstrated to me the outlook and perspective necessary to build a lab from scratch. I also wish to thank all of the grad students I've worked with in our group: Robert Ariniello, Keenan Hunt-Stone, Valentina Lee, and Claire Hansel. It is with their help and guidance I am able to survive the perils of the laser lab and of mountains, alike.

During my time at FACET-II I was able to work with a number of talented individuals and learned so much about how science is performed at large facilities. For this experience I thank Mark Hogan, Christine Clarke, Spencer Gessner, Brendan O'Shea, Doug Storey, Claudio Emma, Elias Gerstmayr, Zhijiang Chen, James Allen, Juan Cruz, and the many additional FACET collaborators across the U.S. and abroad. I owe a lot to Henrik Ekerfelt for taking me under his wing and surviving with me down in the tunnels during late nights shifts.

I couldn't have made it this far without the friendship and encouragement from some incredible people. Laken Adkins, Austin Anderson, Kyle Krowpman, Sam Munnis, Zac Simon, and Josh Vickery are the greatest human beans I could have the privilege of knowing. I also send my gratitude out to Chad Glaze for always providing a second home to return to, to Ethan Clements for carousing with me in the barrel and teaching me to enjoy lifting, and to my DnD buddies Logan Hott, Adam Rison, and Chris Vickery for keeping me sane and creative during quarantine.

Lastly, I want to thank all of the never-ending support from my family throughout.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

The pursuit of scientific knowledge requires the necessary tools to make these discoveries, and for the past century particle accelerators have been an invaluable asset to countless endeavors. In the realm of particle and nuclear physics, accelerators allow for physicists to probe the inner structure of subatomic particles and make new fundamental discoveries, such as the more-recent measurement of the Higgs boson. This high energy physics research is powered through using accelerated particles in colliders to measure high energy collision processes. Particle beams can also be manipulated so that they emit bright light to probe small-scale structures. This can take the form of synchrotron machines or free electron lasers (FEL) that produce bright x-ray light. Accelerators even have applications beyond that of pure physics research, reaching into medical, biological, and industrial applications using the various properties of an accelerated particle beam.

It should come as no surprise then that there is a lot of interest and motivation in improving the performance of these accelerators and making them more accessible. Such progress is accomplished two ways: by increasing the maximum energy and luminosity possible by an accelerator, and by making accelerators more compact to reduce their size and cost. Increasing the maximum energy is primarily achieved by simply making larger accelerators. With conventional radio frequency (RF) devices that use electromagnetic fields inside metallic cavities to accelerate charged particles, accelerating gradients can reach as high as 50 MeV/m before dielectric breakdown begins to damage the cavity itself [1].

With RF breakdown as the upper limit, accelerators need to be built larger and larger to

reach higher energies. One such example is the proposed International Linear Collider (ILC), a TeV-scale electron-positron collider with a required footprint of 31 km [2]. Circular accelerators such as the 27 km circumference, 13 TeV Large Hadron Collider (LHC) at CERN use the circular geometry of its design to accelerate particles over many passes. For these accelerators, instead of accelerating gradient as the main limitation one instead has to contend with energy loss through synchrotron radiation and the maximum strength of bending magnets. Synchrotron radiation scales with the relativistic factor to the power of four, γ_L^4, and so electrons with a mass 2000 times less than heavier protons are much more susceptible to this effect. Typically, lighter electrons and positrons are accelerated with linear accelerators to avoid synchrotron radiation, while protons and heavy ions are accelerated with circular accelerators. Even circular accelerators still need to grow bigger to get to higher energies; the proposed Very Large Hadron Collider (VLHC) is an upgrade to the LHC that would reach 100 TeV using an accelerator circumference of up to 100 km.

While building these large accelerators would achieve their respective scientific goals, the cost of such endeavors is not irrelevant and can make funding such projects difficult. Rather than building larger accelerators, one can instead think of improving the technology and build smaller accelerators with greater electromagnetic field strengths for acceleration and focusing. One such pathway is through the use of plasma. Not only is a plasma capable of producing incredible large electromagnetic fields within its volume, but a plasma is already by definition broken down into electrons and ions so there is no RF breakdown limit. In particular, fields produced by a density wave excited in a plasma using a relativistic particle beam is the basis for plasma wakefield acceleration (PWFA) [3, 4, 5, 6]. This acceleration technique effectively uses the plasma as a medium to transfer energy from an initial drive bunch to a secondary witness bunch, and this transfer can occur over a very short distance.

At SLAC National Accelerator Laboratory, experiments at the Final Focus Test Beam facility were able to demonstrate accelerating gradients of up to 52 GeV in a plasma wake driven by an electron beam [7]. The next step was to introduce a witness beam into the accelerating phase of this plasma wake. This was performed at SLAC's Facility for Advanced Accelerator Experimental

Tests (FACET) [8], and initially demonstrated high energy transfer efficiency but with a low total energy gain by the witness bunch [9]. This would be improved at FACET to demonstrate 9 GeV of energy gain in a single PWFA stage [10], and FACET has since been upgraded to FACET-II in an effort to improve the quality of the accelerated bunch [11]. In parallel, the field of laser wakefield acceleration (LWFA) has been using high-intensity lasers to drive the plasma wake instead of a relativistic electron beam, and typically generating the accelerated bunch within the plasma wake itself. This regime can generate accelerating fields of over 100 GeV/m and electron beams of 4.2 GeV [12, 13] and continues to improve on the total charge and quality of the accelerating beams.

Along with the accelerating gradient in an accelerator, plasma-based devices have the capability of improving upon transverse fields that focus particle beams. Along with the need to reach higher energies, particle beams for colliders also need to be focused to as high density as possible to increase the collision rate for events. Conventional technology of magnetic quadrupoles are capable of focusing beam with magnetic field gradients of $1-100$ T/m, and permanent magnetic quadrupole (PMQ) devices can improve upon these gradients up to 500 T/m [14]. These quadrupole focusing structures also have the limitation where they are capable of focusing in one transverse plane while defocusing in the other plane, which further complicates and limits focusing lattices for charged particles beams. However, just as with the accelerating fields, plasmas are able to facilitate electromagnetic fields strengths that can exceed what is possible in metallic structures by leveraging the high density of charge available in laboratory plasma sources. One such device is the active plasma lens [15], where the focusing on a charge particle beam is due to an induced longitudinal current through the plasma. Active plasma lenses have been characterized providing axisymmetric focusing gradients on the order of several kT/m while preserving the transverse quality of the electron beam [16, 17, 18, 19].

Another regime of plasma lensing is the passive plasma lens. While the active plasma lens uses an applied current in the plasma to focus the particle beams, a passive plasma lens operates essentially as a short PWFA. In addition to the accelerating fields of a plasma wake, a strong

axisymmetric focusing force also exists. The exact nature of the passive plasma lens focusing force is determined by how dense the plasma is with respect to the beam density. In the overdense regime, the plasma is much more dense than the beam and the focusing force scales with the beam density through linear plasma motion [20, 21, 22]. On the other hand, the underdense regime is characterized by an electron beam much more dense than the plasma and the axisymmetric focusing force scales with the plasma density through Coulomb attraction between the negatively charged electron beam and the positively charged plasma ions [23, 24]. Such underdense plasma lenses have been used as a plasma jet focusing an electron beam [25], and as a secondary plasma source used to control the high-divergence electron beams generated within an LWFA [26, 27]. This regime has favorable focusing fields for preserving the transverse quality of an electron beam, while also providing equivalent focusing fields on the order of MT/m, orders of magnitude higher then electromagnetic or permanent magnet focusing devices.

The subject of this dissertation is on improving the design of these passive, underdense plasma lenses. With the large focusing strengths present in these plasma wakes, knowledge and control of the plasma thickness is crucial to determining how much focusing force an electron beam will receive. It is also important to be able to quantify the quality of the electron beam after it has been focused by a passive underdense plasma lens, as this determines the minimum possible spot size the beam can be focused to. Here, we investigate how lasers can be used to precisely and accurately ionize a thin volume of plasma with thickness on the order of $10's$-$100's$ mum, and the resulting effects on an electron beam's 6D phase space.

We also in this dissertation propose an experimental design for performing underdense plasma lens experiments at FACET-II using two 10 GeV electron bunches. This design leverages laser ionization to generate a plasma lens in a straightforward setup that is also capable of adjusting the parameters of the plasma lens to vary the focusing force. We report on progress towards the realization of these experimental goals made during commissioning shifts at FACET-II.

1.1 Outline

Organization of this thesis is as follows. Chapter 2 is a review of the plasma and accelerator physics concepts used in the formalism of the remainder of the dissertation. This chapter derives particle beam motion in external fields and introduces parameters to characterize the beam's quality in phase space. Also, we present the necessary background information to analytically understand the plasma dynamics and electromagnetic fields within a plasma wake.

Chapter 3 narrows the analytic work on applications directly relevant to passive underdense plasma lenses. This chapter contains some additional background work performed previously in describing the focal length and focusing strengths of underdense plasma lenses, but also contains newer research on applications of these plasma lenses towards other current goals in accelerator physics. Such examples include emittance preservation in PWFA where strong and accurate focusing of the beam is needed to limit transverse quality degradation during PWFA acceleration, and studying the fundamental focusing limit on electron beams due to synchrotron radiation. We also cover numerical efforts to simulate underdense plasma lenses in more exotic regimes.

Chapter 4 moves on to our work on designing an underdense plasma lens using laser ionization. We investigate schemes to generate thin plasma lenses of thicknesses $10's$-$100's$ mum in both static volumes of gas and in outflows of supersonic gas jets. Along with these designs, we taken into consideration the impacts of plasma refraction on the ionization process and the impacts of the nonuniform gas density of a gas jet outflow on the plasma lens focusing fields.

Chapter 5 describes our experimental setup at FACET-II, and includes the description of the ionizing laser as well as the electron beam at the experimental area. We introduce the various diagnostic equipment available for the plasma lens experiment, and discuss the progress made towards a fully operational underdense plasma lens experiment at FACET-II.

Chapter 6 covers a recent commissioning shift performed at FACET-II. While the laser and electron beam parameters are still getting tuned up to the final parameters we would like for an optimal experiment, this commissioning shift allowed us to gain useful experience in running the

experiment. We also analyze the data to look for possible underdense plasma lens focusing in this sub-optimal regime using a single, large electron beam to both drive the plasma wake and be partially focused.

Lastly, in Chapter 7 we summarize the dissertation and outline the next steps for the plasma lens experiment at FACET-II and the future outlook of underdense plasma lenses as a whole.

Chapter 2

Electron Beams and Plasma Wakes

Passive plasma lenses and plasma wakefield acceleration utilize the intense electromagnetic fields excited in a plasma when a relativistic electron beam passes through a plasma. These fields are able to both accelerate particles with a high amount of energy gain in a short distance, and are capable of focusing beams with a force stronger than conventional magnets found in linear accelerators. By modeling these fields on a particle beam made up of a distribution of particles we can describe the evolution of the beam's phase space and quantify the quality of the accelerated and/or focused beams. To do so, we must introduce the formalism to two areas of Physics which describe these processes. First, using accelerator physics we come up with a formalism to model how a particle beam evolves in phase space under the influence of external electromagnetic fields. Second, using plasma physics we model how a plasma wave behaves in the linear and nonlinear regimes in find what happens when a particle beam is placed within such a wave.

2.1 Accelerator Physics

We begin our background information overview by introducing the element we are interested in focusing: the relativistic electron beam. Electron beams are a vital asset to many fields of Science, and so there is a wealth of analytic and numerical techniques for modeling these beams. Here, we go through the Courant-Snyder formalism for transverse dynamics of relativistic electron beams, as well as some aspects of the beam to keep in mind as we introduce dynamics with beam-plasma interactions.

2.1.1 Ultra-Relativistic Electron Beams and Accelerators

First, we must answer the question of what exactly is an electron beam. A general definition is that an electron beam is simply a collection of electrons that all have similar positions and momentum. The size of an electron beam can vary depending on the optics an accelerator uses to focus and transport the beam, but at interaction points for collisions these beams can have transverse spot sizes on the order of nanometers. We primarily consider electron beams in this dissertation, but positron beams are also important in the world of particle physics and have similar properties to electron beams.

Electrons themselves are very light particles. To reach speeds of 99% of the speed of light light, an electron only needs to have 3.6 MeV total energy. Modern accelerator-based research facilities such as SLAC operate with electron beams with energies on the order of 10 GeV, and for all practical purposes we can treat propagating electron beams in the ultra-relativistic limit. In this limit we safely approximate the velocity for every particle in the beam as the speed of light, c. Even if the particles have a distribution of energy, their near-equivalent velocities mean that the beam does not spread out longitudinally.

Particle accelerators got their start in 1911 in the wake of Rutherford discovering the atomic nuclei and scatting α-particles off of a gold foil.[1] Since then, machines have been built to effectively provide accelerating electric fields to charged particle beams. Conventional accelerators typically accelerate particles using a radio-frequency (RF) cavity, which is a metallic waveguide typically made of copper for room temperature operation or more exotic materials, such as niobium, for low-temperature, superconducting cavities. These structures allow for a charged particle beam to pass through the central axis, and an electromagnetic RF pulse is sent through the structure at the same time the electron beam passes through. These cavities are designed in such a way that the RF electromagnetic fields within the cavity are spatially uniform and are timed such that the charged particle beam always sees an accelerating phase of the electric field. In each pass of a cavity a beam

[1] If we are stretching the definition of 'accelerator' then, quoting Lee in Ref. [28], "the first accelerator dates back to prehistoric-historic times, when men built bows and arrows for hunting."

gains some amount of energy, with copper allowing for a maximum of 50 MV/m before dielectric breakdown [1].

In addition to these accelerating structures, accelerators also have many other electromagnetic optics in the beamline to provide transverse forces to a particle beam. For instance, quadrupole magnets exert focusing and defocusing forces on a beam, while a dipole magnet is capable of bending a beam's trajectory. Other beamline elements exist for manipulating the phase space of a beam and keeping it contained in the accelerator as it gets accelerated.

The high relativistic Lorentz factors for particle beams also limits how accelerators are built. From Jackson Sec. 14.2 [29] when high energy particle beams travel in a curved trajectory, such as in the field of a dipole magnet, they radiate power at a rate of

$$P = \frac{2}{3}\frac{e^2 c}{\rho^2}\beta^4\gamma^4, \tag{2.1}$$

with $\beta = v/c$, and ρ as the orbit radius.[2] If we consider the 13 TeV Large Hadron Collider of CERN in Europe, the accelerator is built as a large circle with a bending radius capable of handling the energy loss of heavier particles. For an electron beam however, the energy loss in a curved trajectory is much larger and so it is necessary to build straight linear accelerators to reach the highest energies of interest to the particle physics community. An example of this is the 3 km electron accelerator at SLAC. Synchrotron radiation of Eqn. 2.1 can be purposefully generated by electron beams as a light source by sending relativistic electron beams through an undulator of alternating dipole magnets, such as the LCLS facility at SLAC [30]. It is also possible for strong focusing optics to generate synchrotron radiation from the focusing or defocusing of electrons even if the beam as a whole is centered in the focusing optic.

2.1.2 Transverse Beam Dynamics

To begin modeling an electron beam, we first formulate how to model how a single relativistic electron interacts with applied electromagnetic fields. Here I will be using the formalism set out by

[2] A result first obtained by Alfred-Marie Liénard in 1898.

Lee in Ref. [28]. Let us assume that we have an electron propagating along a reference orbit $r_0(s)$ with s as the propagation direction with a relativistic speed determined by its Lorentz factor, γ_L. The reference orbit $r_0(s)$ can exist in a plane, and so $\hat{x}$ is defined to be perpendicular to s and in the tangent plane while $\hat{y}$ is defined to be perpendicular to both s and the tangent plane. In other words, $\hat{x}$ is typically the horizontal axis and $\hat{y}$ is typically the vertical axis. This is known as the Frenet-Serret coordinate system:

$$\hat{s}(s) = \frac{d\vec{r}_0(s)}{ds}, \quad \hat{x}(s) = -\rho(s)\frac{d\hat{s}(s)}{ds}, \quad \hat{y}(s) = \hat{x}(s) \times \hat{s}(s) \tag{2.2}$$

with $\rho(s) = \left|\frac{ds}{d\hat{s}}\right|$ as the radius of curvature. For most purposes in this thesis, we replace these general coordinates with the simpler definition where $r_0(s)$ to be a straight line in the $\hat{z}$ direction.

The particle's position is given in Cartesian coordinates as

$$\vec{\boldsymbol{x}}(t) = x(t)\hat{x} + y(t)\hat{y} + z(t)\hat{z}. \tag{2.3}$$

However, it is more useful to express Eqn. 2.3 in terms of the longitudinal position s in the reference orbit. Here we make the substitution

$$\xi = v_0 t - s \tag{2.4}$$

and Eqn. 2.3 is now written

$$\vec{\boldsymbol{x}}(s) = x(s)\hat{x} + y(s)\hat{y} + \xi(s)\hat{z}. \tag{2.5}$$

The reference particle velocity v_0 is given by particle mass m and momentum p_0 as $v_0 = \sqrt{\frac{c^2 p_0^2}{m^2c^2+p_0^2}}$, and for relativistic electron beams we approximate $v_0 \approx c$. In the coordinate system of Eqn. 2.5 a particle's 6D phase space is fully defined with $(x, p_x, y, p_y, \xi, p_s)$. For convenience, we instead work in trace space coordinates where the transverse momenta p_x and p_y are instead represented by the angular directions

$$x' = \frac{dx}{ds} = \frac{p_x}{p_s}; \quad y' = \frac{dy}{ds} = \frac{p_y}{p_s}. \tag{2.6}$$

We also express the longitudinal momentum as a fractional deviation from the reference particle, $\delta = (p - p_0)/p_0$. With this new notation, the 6D phase space of a particle is defined through $(x, x', y, y', \xi, \delta)$.

To model the evolution of a particle's transverse phase space under the influence of external electromagnetic fields, we begin with the Lorentz force law:

$$\frac{d\vec{\boldsymbol{p}}}{dt} = \vec{\boldsymbol{F}} = q(\vec{\boldsymbol{E}} + \vec{\boldsymbol{v}} \times \vec{\boldsymbol{B}}). \tag{2.7}$$

Under the ultra-relativistic approximation, $v \approx c$ and $p \approx \gamma_L mc$. Looking at only the x component the Lorentz force, the equation of motion for a particle of charge q becomes

$$\frac{dp_x}{ds} = \frac{q}{c}E_x + q\left(\frac{v_y}{c}B_z - \frac{v_z}{c}B_y\right). \tag{2.8}$$

Using the definitions of the trace space coordinates of Eqn. 2.6 and assuming these angles are small ($x' \ll 1$), the equation of motion is simplified to

$$x'' + \frac{\gamma_L'}{\gamma_L}x' = \frac{q}{\gamma_L mc^2}E_x + \frac{q}{\gamma_L mc}(x'B_z - B_y). \tag{2.9}$$

For the case where $B_z = 0$, the x-y coupling is removed from Eqn. 2.9. This is true for many scenarios involving magnetic accelerator optics and beam-plasma interactions. We also can include both the external electric and magnetic fields in a plasma wakefield $W_x \equiv E_x - cB_y$. For simplicity, we also now assume here that the electron is neither losing or gaining energy so $\gamma_L' = 0$. We now have

$$x'' = \frac{q}{\gamma_L mc^2}W_x x, \tag{2.10}$$

with a similar equation for the y coordinate. To go further, we need to consider what the external electromagnetic fields are. A good example to work off of is that of a quadrupole magnet (Fig. 2.1). These provide magnetic fields that are orthogonal to $\hat{s}$ and have strengths proportional to the distance from reference orbit, or $B_x = K_x x$. This allows us to write

$$x'' + K_x(s)x = 0, \tag{2.11}$$

which is referred to as Hill's equation. The variable K is referred to as the focusing strength and has units of m^{-2}. This equation is a linear, homogeneous, second order differential equation that

has different solutions depending on the sign of $K_x(s)$:

$$x(s) = \begin{cases} a\cos\left(\sqrt{K}s + b\right), & K > 0 \\ as + b, & K = 0 \\ a\cosh\left(\sqrt{-K}s + b\right), & K < 0 \end{cases} \tag{2.12}$$

A $K > 0$ indicates a focusing force on the particle, $K < 0$ a defocusing force, and $K = 0$ is a drift space without an external field. This is important when considering quadrupole magnets because, as seen in Fig. 2.1, quadrupoles have the signs flipped between the two transverse directions, or $K_x = -K_y$.

The solution of Eqn. 2.12 can be further interpreted in the formalism of transfer matrices to replace the integration constants a and b with the initial conditions $x(s_0)$ and $x'(s_0)$:

$$\boldsymbol{x}(s_0) = \begin{pmatrix} x(s_0) \\ x'(s_0) \end{pmatrix}; \quad \boldsymbol{x}(s) = M(s|s_0)\boldsymbol{x}(s_0). \tag{2.13}$$

The transfer matrices $M(s|s_0)$ for the three solutions of Eqn. 2.12 are

$$M(s|s_0) = \begin{cases} \begin{pmatrix} \cos\sqrt{K}L & \frac{1}{\sqrt{K}}\sin\sqrt{K}L \\ -\sqrt{K}\sin\sqrt{K}L & \cos\sqrt{K}L \end{pmatrix} & K > 0 \\ \begin{pmatrix} 1 & L \\ 0 & 1 \end{pmatrix} & K = 0 \\ \begin{pmatrix} \cosh\sqrt{|K|}L & \frac{1}{\sqrt{|K|}}\sinh\sqrt{|K|}L \\ \sqrt{|K|}\sinh\sqrt{|K|}L & \cosh\sqrt{|K|}L \end{pmatrix} & K < 0 \end{cases} \tag{2.14}$$

where $L = s - s_0$ is the length of the matrix element. An example of how a focusing quadrupole and drift space effects the phase space of an electron beam is shown in Fig. 2.2. The solutions for a nonzero K are valid in modeling the transverse dynamics of an electron in other focusing optics, as we will see in Sec. 2.3.3 for the case of the transverse wakefield inside of a plasma wake.

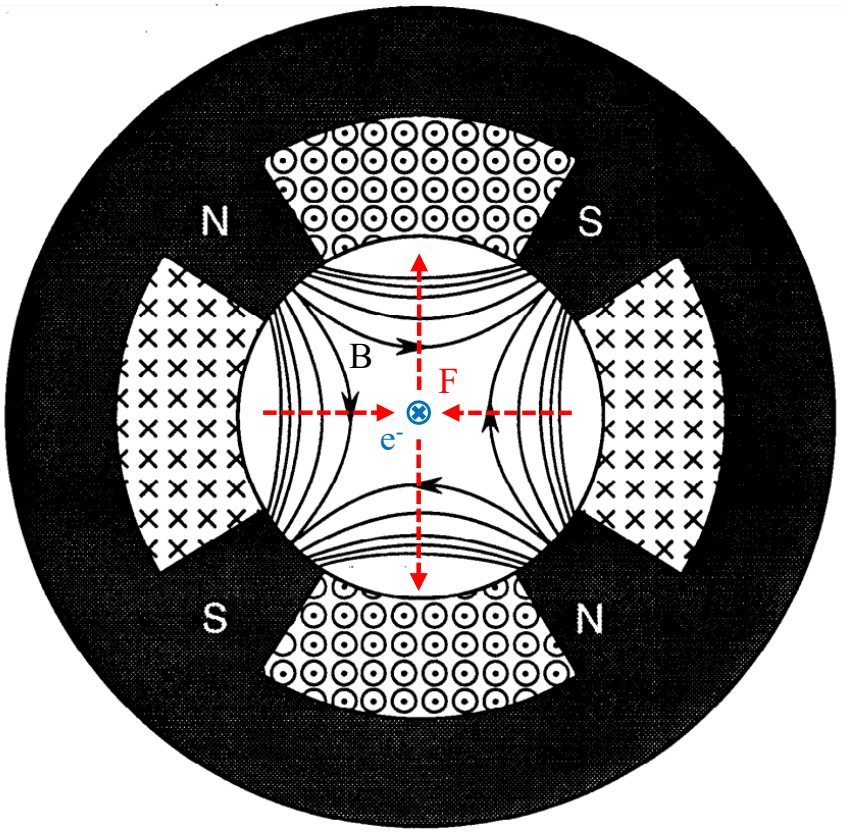

Figure 2.1: Layout of a quadrupole electromagnet. The magnetic fields are plotted in solid black, and their strengths increase approximately linear from the central axis. For an electron beam propagating into the page (blue), the direction of the transverse force is shown in red. Quadrupoles focus charged particle beams in one transverse direction and defocus in the other. Adapted from Ref. [31]

For quadrupole magnets that are relatively short, $\sqrt{K}L \ll 1$, we can approximate these transfer matrices to the thin regime:[3]

$$M_{thin} = \begin{pmatrix} 1 & 0 \\ \pm 1/f & 1 \end{pmatrix}. \tag{2.15}$$

Here $f = \frac{1}{|K|L}$ is the focal length of such a thin lens quadrupole where only the particle's transverse momentum is affected and its transverse position remains unchanged. For a focusing quadrupole, the sign in Eqn. 2.15 is negative.

[3] While most quadrupole magnets are on the order of a meter in length, K can be low enough for the thin lens regime to be valid.

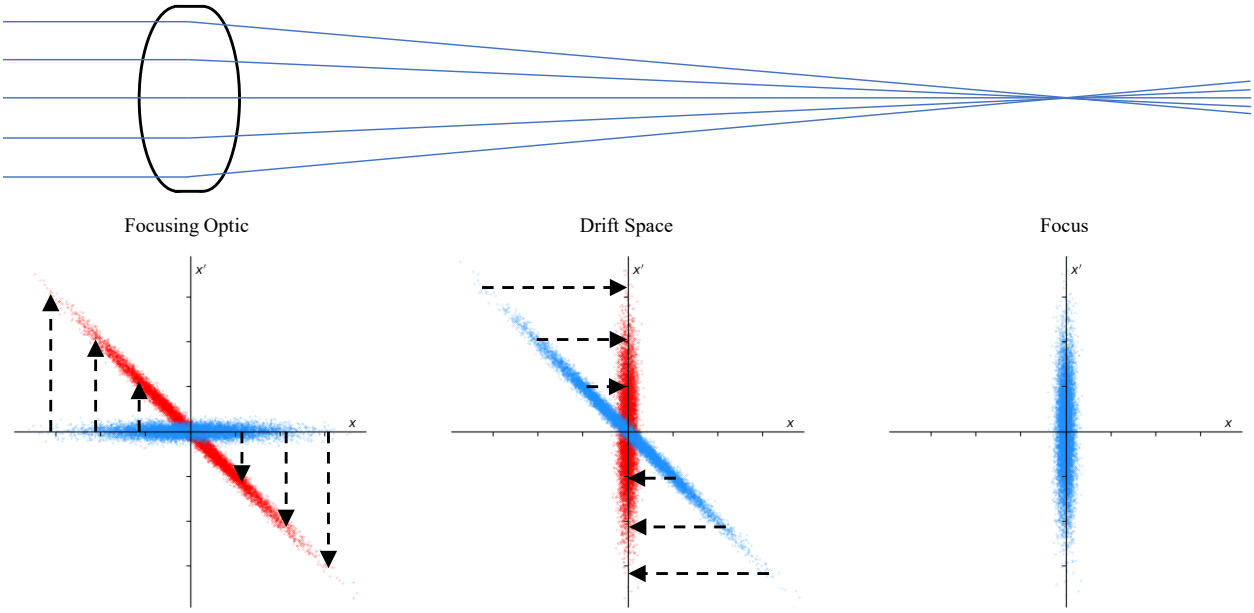

Figure 2.2: Evolution of an electron beam's phase space as it passes through a focusing quadrupole and into subsequent drift space (top). The three subplots on the bottom show a distribution of particles, with blue as the phase space before and red as the phase space after. In the focusing optic, the particle positions are unchanged but its momentum changes according to its position. In the drift space, the momentum is unchanged while the positions shift relative to its momentum. At the focus the size of the beam is at a minimum.

2.1.3 Courant-Snyder Parameters for a Single Particle

While the general solution using the phase space coordinates x and x' is valid, it is more useful to define the solution of Hill's equation in parameters more applicable to the form of the solution. These Courant-Snyder (CS) parameters will give more insight to the Hill's equation solution and simplify the modeling of transverse beam dynamics. We start here by defining these parameters with the Hill's equation solution for a single particle's transverse motion, and in Sec. 2.1.4 we redefine these CS parameters using statistical distributions of a beam with many particles.

Following the work of Courant and Snyder [32] with additional information from Lee [28], if we return to Hill's equation (Eqn. 2.11) we can use the general solution:

$$x(s) = ax_1(s) + bx_2(s), \tag{2.16}$$

$$x_1(s) = \omega(s)e^{i\Phi(s)}, \tag{2.17}$$

$$x_2(s) = \omega(s)e^{-i\Phi(s)}, \tag{2.18}$$

with $\omega(s)$ and $\Phi(s)$ as arbitrary functions for the amplitude and phase with longitudinal dependence on s. If we substitute Eqn. 2.17 back into Hill's equation we arrive at the following constraints for $\omega(s)$ and $\Phi(s)$:

$$\omega'' + K(s)\omega - \frac{1}{\omega^3} = 0, \qquad \Phi' = \frac{1}{\omega^2}. \tag{2.19}$$

Similarly, the solution $x_2(s) = \omega(s)e^{-i\Phi(s)}$ also satisfies Hill's equation and is independent from x_1.

We can evaluate Eqn. 2.19 by defining the following parameters:

$$\beta = \omega^2, \tag{2.20}$$

$$\alpha = -\frac{\beta'}{2}, \tag{2.21}$$

$$\gamma = \frac{1+\alpha^2}{\beta}. \tag{2.22}$$

The parameters α, β, and γ are referred to as Courant-Snyder (CS) parameters.[4] These definitions for the CS parameters are useful when evaluating the trajectory of a single particle, but we will use a different definition when describing the full beam distribution. (See Eqn. 2.47 for the statistical definitions) Using β in Eqn. 2.19, we find a differential equation for β

$$\frac{1}{2}\beta'' + K\beta - \frac{1}{\beta}\left[1 + \left(\frac{1}{2}\beta'\right)^2\right] = 0 \tag{2.23}$$

that is evaluated to find the evolution of β for the single particle. We can enforce a periodic boundary condition on β such that $\beta(s) = \beta(s+l)$ and $\beta'(s) = \beta'(s+l)$ for a K that is also assumed to be periodic, where l is the same period as for a periodic K. The general solution to Hill's equation comes from finding $x_1(s)$ and $x_2(s)$, and then applying initial conditions $x_0 = x(s_0)$ and $x_0' = x'(s_0)$ to find $x(s) = ax_1(s) + bx_2(s)$.

Returning to the single particle CS parameters, we can uncover additional motivation for the definitions of β, α, and γ by writing the transfer matrix solution to Hill's equation with the CS

[4] Not to be confused with the β_L and γ_L used for relativistic longitudinal speed and energy of the beam.

parameters:

$$M(s_2|s_1) = \begin{pmatrix} \sqrt{\frac{\beta_2}{\beta_1}}(\cos\Phi + \alpha_1\sin\Phi) & \sqrt{\beta_1\beta_2}\sin\Phi \\ -\frac{1+\alpha_1\alpha_2}{\sqrt{\beta_1\beta_2}}\sin\Phi + \frac{\alpha_1-\alpha_2}{\sqrt{\beta_1\beta_2}}\cos\Phi & \sqrt{\frac{\beta_1}{\beta_2}}(\cos\Phi - \alpha_2\sin\Phi) \end{pmatrix}. \tag{2.24}$$

Next, enforcing the periodic boundary conditions results in equivalent CS parameters at s_1 and s_2. Using our definition of γ we simplify Eqn. 2.24 down to

$$M(s+l|s) = \begin{pmatrix} \cos\Phi + \alpha\sin\Phi & \beta\sin\Phi \\ -\gamma\sin\Phi & \cos\Phi - \alpha\sin\Phi \end{pmatrix} = \boldsymbol{I}\cos\Phi + \boldsymbol{J}\sin\Phi, \tag{2.25}$$

where $\boldsymbol{I}$ is the identity matrix and $\boldsymbol{J}$ includes the CS parameters

$$\boldsymbol{J} = \begin{pmatrix} \alpha & \beta \\ -\gamma & -\alpha \end{pmatrix}. \tag{2.26}$$

For the solution to Hill's linear equation to be valid, $\boldsymbol{J}$ must have a determinate of one so we have

$$\beta\gamma - \alpha^2 = 1. \tag{2.27}$$

This is equivalent to the definition of γ in Eqn. 2.22. Because of this relation, one only needs to know two of the three CS parameters to be able to calculate the third.

We can then write the motion of a single particle using the linear combination of solutions in Eqn. 2.16 with an initial phase and amplitude as

$$x(s) = \sqrt{\epsilon\beta(s)}\cos(\Phi(s) + \Phi_0), \tag{2.28}$$

where $\Phi(s)$ is the betatron phase advance and is found by integrating Eqn. 2.19:

$$\Phi(s) = \int_0^s \frac{ds'}{\beta(s')}. \tag{2.29}$$

The initial amplitude ϵ is the single particle emittance, and it is an invariant of the CS parameters:

$$\epsilon = \gamma x^2 + 2\alpha x x' + \beta x'^2. \tag{2.30}$$

This single particle emittance can be plotted as an ellipse in phase space, as shown in Fig. 2.3. As a particle travels down an accelerator, its transverse position and angle will be located on this

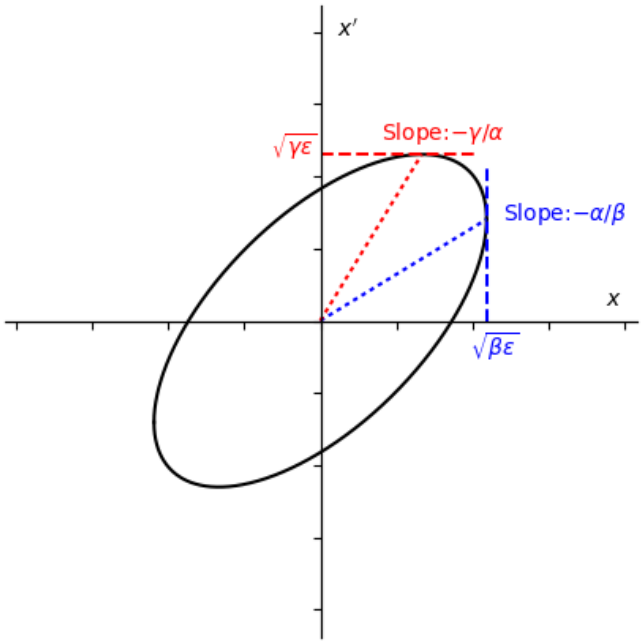

Figure 2.3: Phase space ellipse of the single particle emittance, ϵ, which is calculated as the area of the ellipse. The shape of the ellipse is determined by the CS parameters and the position of the particle on this ellipse is determined by the betatron phase advance.

emittance ellipse at a position determined by the betatron phase advance. If the reference orbit is defined as a particle existing at $x = x' = 0$, then the single-particle emittance is a measure of the deviation from the reference orbit.

It is also possible to write the evolution of the CS parameters as a transfer matrix. If the phase space transfer matrix M for x and x' is known, then we can use the following to propagate CS parameters from longitudinal position s_1 to s_2:

$$\begin{pmatrix} \beta_2 \\ \alpha_2 \\ \gamma_2 \end{pmatrix} = \begin{pmatrix} M_{11}^2 & -2M_{11}M_{12} & M_{12}^2 \\ -M_{11}M_{21} & M_{11}M_{22} + M_{12}M_{21} & -M_{12}M_{22} \\ M_{21}^2 & -2M_{21}M_{22} & M_{22}^2 \end{pmatrix} \begin{pmatrix} \beta_1 \\ \alpha_1 \\ \gamma_1 \end{pmatrix}. \tag{2.31}$$

Here, the values M_{ij} represent the individual elements in the transfer matrix M.

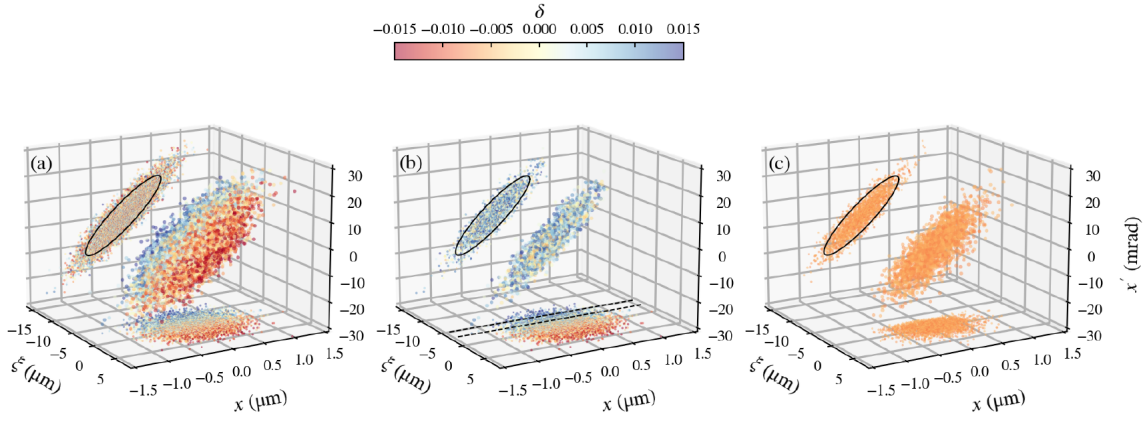

Figure 2.4: 4D phase space of a particle beam. Each dot represents a single electron, and its color represents its relative energy offset δ from the centroid energy value. The "back-wall" of subplot (a) shows the x-x' projection for the full beam, subplot (b) shows the projection for a single longitudinal slice, and subplot (c) shows the projection for a single energy slice. Reproduced from Ref. [33]

2.1.4 Transverse Beam Distribution and Beam Emittance

As mentioned previously in Sec. 2.1.1, an electron beam is comprised of a distribution of electrons. If we wish to model the evolution of the electron beam as a whole, it is much easier to come up with a formalism for describing the beam's distribution rather than track each electron individually. Here, we perform the calculation which led to the single particle emittance for a full particle beam distribution and find a geometric emittance for the beam. This geometric emittance is related to statistical moments of the beam's distribution, and we use these statistical moments to further define Courant-Snyder parameters that relate to the full beam distribution rather than a single particle.

A starting point is the full 6D phase space of of an electron beam, $f(x, x', y, y', \xi, \delta)$. As we saw with Hill's equation, typically the x and y planes behave independently of one another so we can limit this distribution to only the 4D distribution in x and ξ, or $f(x, x', \xi, \delta)$ and assume that the respective y subspace will yield similar results.

From this 4D distribution we can take several projections. First is a projection onto the x-x'

plane that is gives the full beam's phase space in x:

$$f_\perp(x,x') = \int d\xi d\delta \; f(x,x',\xi,\delta). \tag{2.32}$$

We can also project single, longitudinal slices of ξ onto the x-x' phase space plane,

$$f_\xi(x,x';\xi) = \frac{1}{N(\xi)} \int d\xi' d\delta \; f(x,x',\xi,\delta)\delta_D(\xi'-\xi), \tag{2.33}$$

as well as single slices of a particular energy δ,

$$f_\delta(x,x';\delta) = \frac{1}{N(\delta)} \int d\xi d\delta' \; f(x,x',\xi,\delta)\delta_D(\delta'-\delta). \tag{2.34}$$

Here N is the number of particles in that particular slice, $N(\xi) = \int dxdx'd\delta f(x,x',\xi,\delta)$ and $N(\delta) = \int dxdx'd\xi f(x,x',\xi,\delta)$, and δ_D is the Dirac delta function. Projections of such a 4D phase space are shown in Fig. 2.4. For all three of these projections, we can take moments of the distribution in the x-x' plane to better characterize them:

$$\begin{aligned}
\langle x\rangle &= \int dxdx' x f_\perp(x,x'), \qquad \langle x'\rangle = \int dxdx' x' f_\perp(x,x'),\\
\langle x^2\rangle &= \int dxdx' x^2 f_\perp(x,x'), \quad \langle x'^2\rangle = \int dxdx' x'^2 f_\perp(x,x'),\\
\langle xx'\rangle &= \int dxdx' xx' f_\perp(x,x').
\end{aligned} \tag{2.35}$$

These moments can be used to calculate RMS spot sizes and divergences:

$$\sigma_x^2 = \langle x^2\rangle - \langle x\rangle^2, \;\; \sigma_{x'}^2 = \langle x'^2\rangle - \langle x'\rangle^2, \;\; \sigma_{xx'} = \langle xx'\rangle - \langle x\rangle\langle x'\rangle. \tag{2.36}$$

We can use this formalism to statistically describe for an electron beam of arbitrary distribution, although in practice electron beams commonly have Gaussian distributions.[5] For a Gaussian distribution, these statistical RMS spot sizes and divergences are equivalent to the analytic Gaussian spot sizes for the distribution.

If we now return to the definition for single particle emittance (Eqn. 2.30) we can consider what an emittance for a full electron beam would be. First, let's assume that the beam we are

[5] For at least the transverse components. The longitudinal components can often be non-Gaussian, with the energy distribution commonly being close to a flattop distribution.

characterizing has every particle with equivalent energy. This would also be the same as using an energy slice of the beam from Eqn. 2.34. For such a distribution, the transfer matrix that describes the particle evolution is the same for every particle in the beam. For the entire beam,

$$X(s) = \begin{pmatrix} \boldsymbol{x}_1 & \boldsymbol{x}_2 & \cdots & \boldsymbol{x}_N \end{pmatrix} = M(s|s_0)X(s_0), \tag{2.37}$$

where X is a $4 \times N$ matrix for the transverse coordinates of each particle. If we assume that the beam is overall centered on x and x' such that $\langle x \rangle = \langle x' \rangle = 0$, then the other beam moments can be found through taking the outer product of the matrix:

$$\Sigma = \begin{pmatrix} \langle x^2 \rangle & \langle xx' \rangle \\ \langle xx' \rangle & \langle x'^2 \rangle \end{pmatrix} = \frac{XX^T}{N}. \tag{2.38}$$

This Σ matrix evolves according to the transfer matrix M of the accelerator,

$$\Sigma(s) = \frac{MX(s_0)[MX(s_0)]^T}{N} = \frac{MX(s_0)X(s_0)X^T(s_0)M^T}{N} = M\Sigma(s_0)M^T, \tag{2.39}$$

and because the determinant of the transfer matrix M is 1, the determinant of Σ remains constant. Using this determinant, we define the geometric emittance

$$\epsilon = \sqrt{\det(\Sigma)} = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2} \tag{2.40}$$

as a constant of motion for the beam that quantifies the beam's transverse deviation from the reference orbit. An example of the geometric emittance for a particle beam is plotted in Fig. 2.5.

Next, we want to relate this geometric emittance to the CS parameters to get a statistical definition for α, β, and γ. Following the derivation of Buon [34], we can rotate the beam distribution by an angle θ to a coordinate system (ω, ω') such that $\sigma_{\omega'}$ is minimized. This gives us the condition

$$\frac{\partial}{\partial \theta} \sigma_{\omega'}^2 = 0. \tag{2.41}$$

A particle's position in this coordinate system is

$$\omega' = x' \cos\theta - x \sin\theta, \tag{2.42}$$

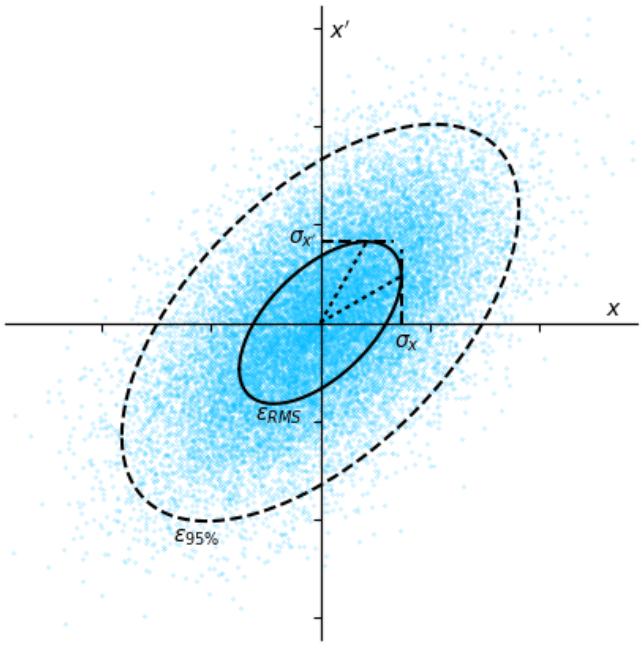

Figure 2.5: Phase space projection of a particle beam with its geometric emittance over-plotted on the distribution. The geometric emittance is related to the RMS size of the beam, σ_x. Alternatively, one could also define an emittance $\epsilon_{95\%}$ which contains 95% of the beam particles within the ellipse.

and using Eqn. 2.41 with Eqn. 2.42 lets us solve for θ:

$$\tan 2\theta = \frac{2\langle xx' \rangle}{\langle x^2 \rangle - \langle x'^2 \rangle}. \tag{2.43}$$

The two moments are evaluated as

$$\sigma^2_{\omega,\omega'} = \frac{1}{2}\left(\langle x^2 \rangle + \langle x'^2 \rangle \pm \frac{2\langle xx' \rangle}{\sin 2\theta} \right), \tag{2.44}$$

and in this coordinate system the geometric emittance is $\epsilon = \sigma_\omega \sigma_{\omega'}$. The ellipse can be defined

$$\frac{\omega^2}{\sigma^2_\omega} + \frac{\omega'^2}{\sigma^2_{\omega'}} = 1. \tag{2.45}$$

Transforming Eqn. 2.45 back to the (x, x') coordinate system by rotating away from (ω, ω'), we find the following for the geometric emittance:

$$x^2\langle x'^2 \rangle - 2xx'\langle xx' \rangle + x'^2\langle x^2 \rangle = \epsilon. \tag{2.46}$$

By comparing this to Eqn. 2.30, we can define statistical CS parameters that correspond to the moments of the beam distribution and to the geometric emittance:

$$\beta = \frac{\langle x^2 \rangle}{\epsilon}, \quad \gamma = \frac{\langle x'^2 \rangle}{\epsilon}, \quad \alpha = -\frac{\langle x x' \rangle}{\epsilon}. \tag{2.47}$$

If the beam is centered on $\langle x \rangle = \langle x' \rangle = 0$, then $\langle x^2 \rangle = \sigma_x^2$, $\langle x'^2 \rangle = \sigma_{x'}^2$, and $\langle x x' \rangle = \sigma_{xx'}$. Taking Fig. 2.5 for an example beam distribution, β corresponds to the physical size of the beam, γ corresponds to the momentum spread, and α is the skew at which this beam ellipse is orientated. Returning to the transfer matrix for every particle in the beam, the Σ matrix of Eqn. 2.39 can be rewritten using CS parameters as

$$\Sigma = \epsilon \begin{pmatrix} \beta & -\alpha \\ -\alpha & \gamma \end{pmatrix}. \tag{2.48}$$

This tells us that the statistical definition of the CS parameters in Eqn. 2.47 allows us to describe the transverse beam dynamics for the full beam. In general, if the initial conditions of the beam are known and can be used to define the CS parameters of Eqn. 2.47, then we can propagate the betafunction β of the beam through an accelerator using Eqn. 2.23 or by using the corresponding transfer matrix and Eqn. 2.31.

Emittance is a crucial quantity in describing electron beams as it is related to the beam size σ_x at interaction points, and in many experiments having the smallest possible beam size is important for having high luminosity. If a collider is built such that all of the focusing magnets are set up to focus an electron beam down to a focus of β^*, then the spot size of the electron beam at the focus is $\sigma_x^2 = \beta^* \epsilon$. While β^* is limited by the magnet strength and lattice design, ϵ is limited by the electron beam quality itself.

2.1.5 Including Energy Gain and Energy Spread

The CS parameters α, β, and γ defined using the formalism in Sec. 2.1.4 are valid for mono-energetic beams that neither gain or lose energy. For realistic beams with a finite energy distribution that undergo acceleration in an accelerator, this definition is not all that useful. To include energy

spread in CS parameters, we can consider using the definitions for the statistical CS parameters of Eqn. 2.47 for every energy slice δ of the beam (Eqn. 2.34). If we do this, then the CS parameters become functions of δ. We could also take the simpler approach and just use the statistical averages across the full beam (Eqn. 2.32). We lose information on how the different energy slices are shaped in phase space, but if the energy spread is low enough then this approximation can be fine. For the remainder of this thesis, when the CS parameters are used by themselves (i.e. α, β, γ) we are assuming the statistical average for the full beam's projection. When the CS parameters are functions of δ (i.e. $\alpha(\delta)$, $\beta(\delta)$, $\gamma(\delta)$) we are using the energy slice interpretation.

To account for energy spread, we must return to our definition for the geometric emittance. The geometric emittance is defined in phase space using the (x, x') coordinate system, where the angle $x' = p_x/p_z$ is not a true representation of momentum. If we suppose that an particle gains some amount of energy, then momentum p_z increases which results in x' decreasing and geometric emittance ϵ also decreasing. This process of adiabatic dampening renders the geometric emittance as a non-conserved quantity for particles that undergo acceleration, which isn't very useful in practical terms.

Rather, what is conserved is the emittance when it is defined through the projection onto the x-p_x plane. Using this projection, we find an emittance of

$$\epsilon_N = \frac{1}{mc}\sqrt{\sigma_x^2\sigma_{p_x}^2 - \sigma_{xp_x}^2}. \tag{2.49}$$

This ϵ_N is what is referred to as "normalized emittance," and it is a much better indicator of the transverse quality of a particle beam. The normalized emittance is conserved through acceleration stages through the replacement of x' with p_x. As normalized emittance is much more important than geometric emittance, when the term "emittance" is used by itself it is understood that the normalized emittance is what is being referred to.

Normalized emittance is related to geometric emittance through the relativistic Lorentz factor,

$$\epsilon_N = \gamma_L\epsilon, \tag{2.50}$$

and so we can use normalized emittance with all of our definitions of the CS parameters by including this factor of γ_L. Thus, the beam sizes are related to the CS parameters and normalized emittance by rewriting Eqn. 2.47 as:

$$\beta = \frac{\gamma_L \sigma_x^2}{\epsilon_N}, \quad \gamma = \frac{\gamma_L \sigma_{x'}^2}{\epsilon_N}, \quad \alpha = -\frac{\gamma_L \sigma_{xx'}}{\epsilon_N}. \tag{2.51}$$

This is assuming that the beam is centered on-axis such that $\langle x \rangle = \langle x' \rangle = 0$, and also similar expressions exist for CS parameters in the y-y' plane.

2.1.6 Beam Self Fields

An important consequence of ultra-relativistic beams is the fields these particle beams generate. Not only can they produce strong electromagnetic fields, but these fields can also interact with the electron beam itself. To describe these fields we need the inhomogeneous wave equation [29], and to derive this we first begin with Maxwell's equations:

$$\nabla \cdot \boldsymbol{E} = \frac{\rho}{\epsilon_0} \tag{2.52}$$

$$\nabla \cdot \boldsymbol{B} = 0 \tag{2.53}$$

$$\nabla \times \boldsymbol{E} = -\frac{\partial \boldsymbol{B}}{\partial t} \tag{2.54}$$

$$\nabla \times \boldsymbol{B} = \mu_0 \boldsymbol{J} + \frac{1}{c^2} \frac{\partial \boldsymbol{E}}{\partial t}. \tag{2.55}$$

We take the curl of Eqn. 2.54 to get

$$\nabla \times (\nabla \times \boldsymbol{E}) = -\frac{\partial}{\partial t} \nabla \times \boldsymbol{B} \tag{2.56}$$

and use Eqn. 2.55 to get

$$\nabla \times (\nabla \times \boldsymbol{E}) = -\mu_0 \frac{\partial \boldsymbol{J}}{\partial t} - \frac{1}{c^2} \frac{\partial^2 \boldsymbol{E}}{\partial t^2}. \tag{2.57}$$

Using a vector identity the left side of this equation can be rewritten as

$$\nabla(\nabla \cdot \boldsymbol{E}) - \nabla^2 \boldsymbol{E} = -\mu_0 \frac{\partial \boldsymbol{J}}{\partial t} - \frac{1}{c^2} \frac{\partial^2 \boldsymbol{E}}{\partial t^2}. \tag{2.58}$$

Lastly, use Eqn. 2.52 to substitute $\nabla \cdot E$ for

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)\boldsymbol{E} = \frac{1}{\epsilon_0}\nabla\rho + \mu_0\frac{\partial \boldsymbol{J}}{\partial t}, \tag{2.59}$$

which is the inhomogeneous wave equation. A similar expression for the magnetic field can be derived by starting with Eqn. 2.55 and performing similar substitutions with Maxwell's equations. The inhomogeneous wave equation for the magnetic field is

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)\boldsymbol{B} = \mu_0\nabla \times \boldsymbol{J}. \tag{2.60}$$

For an electron beam the source terms are $\rho = -en_b$ and $\boldsymbol{J} = -ecn_b\hat{z}$, which we substitute into the inhomogeneous wave equations:

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)\boldsymbol{E} = -\frac{e}{\epsilon_0}\nabla n_b - \frac{e}{\epsilon_0 c^2}c\hat{z}\frac{\partial n_b}{\partial t} \tag{2.61}$$

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)\boldsymbol{B} = -\mu_0 ec\nabla \times (n_b\hat{z}). \tag{2.62}$$

Next, we change to the co-moving coordinate system in the ultra-relativistic limit, $\xi = ct - z$ (Eqn. 2.4). Doing so results in the $\hat{z}$ component of $\nabla^2\boldsymbol{E}$ cancelling out with the $\frac{1}{c^2}\frac{\partial^2 \boldsymbol{E}}{\partial t^2}$ term, and the same for the $\boldsymbol{B}$ equation. The beam density is also constant in this frame such that $\partial n_b/\partial t = 0$. With this, we have

$$\nabla_\perp^2\boldsymbol{E} = -\frac{e}{\epsilon_0}\nabla_\perp n_b \tag{2.63}$$

$$\nabla_\perp^2\boldsymbol{B} = -\mu_0 ec\left(\frac{\partial n_b}{\partial y}\hat{x} - \frac{\partial n_b}{\partial x}\hat{y}\right). \tag{2.64}$$

Now, we take the cross product of $\boldsymbol{v}/c^2$ with Eqn. 2.63 for

$$\nabla_\perp^2\left(\frac{1}{c^2}\boldsymbol{v} \times \boldsymbol{E}\right) = -\mu_0 ec\left(\frac{\partial n_b}{\partial y}\hat{x} - \frac{\partial n_b}{\partial x}\hat{y}\right) \tag{2.65}$$

and substituting Eqn. 2.64 in here gives

$$\boldsymbol{B} = \frac{1}{c^2}\boldsymbol{v} \times \boldsymbol{E}. \tag{2.66}$$

This result allows us to find the magnetic field given the electric field, and is true for any "cold"[6] beam with little-to-no momentum spread (Jackson Sec. 11.10 [29]). We can already see that B_z is

[6] An ultra-relativistic beam's longitudinal momentum is large enough compared to any transverse motion so that the cold beam assumption is valid.

zero, and for the other fields we need to use Gauss's law to calculate the electric field:

$$E_r(r,\xi) = -\frac{e}{\epsilon_0}\frac{1}{r}\int_0^r r'dr'n_b(r',\xi). \tag{2.67}$$

The E_z field is only zero if $v = c$ [35]. For beams with finite energy spread, we can look back at Faraday's law (Eqn. 2.54) and rewrite it as

$$\left(\frac{\partial E_r}{\partial z} - \frac{\partial E_z}{\partial r}\right) = -\frac{\partial B_\phi}{\partial t}. \tag{2.68}$$

Using our result from Eqn. 2.66 and shifting to a co-moving frame of velocity v_b we get

$$\left(-\frac{\partial E_r}{\partial \xi} - \frac{\partial E_z}{\partial r}\right) = -v_b\frac{\partial B_\phi}{\partial \xi} = -\frac{v_b^2}{c^2}\frac{\partial E_r}{\partial \xi}. \tag{2.69}$$

Lastly, using the definition of $\gamma_L = [1-(v/c)^2]^{-1}$ we get the longitudinal electric field:

$$\frac{\partial E_z}{\partial r} = -\frac{1}{\gamma_L^2}\frac{\partial E_r}{\partial \xi}. \tag{2.70}$$

Going back to the radial field E_r, Eqn. 2.67 can be solved if we assume a density profile. If we assume the beam has a typical Gaussian profile in all three dimensions and total charge Q, the charge density can be written

$$\rho(x,y,\xi) = \frac{Q}{(2\pi)^{3/2}\sigma_x\sigma_y\sigma_\xi}\exp\left(-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2} - \frac{\xi^2}{2\sigma_\xi^2}\right). \tag{2.71}$$

Solving Eqn. 2.67 results in the following for the radial electric field:

$$E_r(r,\xi) = \frac{Q}{(2\pi)^{3/2}\sigma_\xi\epsilon_0}\exp\left(-\frac{\xi^2}{2\sigma_\xi^2}\right)\frac{1}{r}\left[1-\exp\left(-\frac{r^2}{2\sigma_r^2}\right)\right]. \tag{2.72}$$

From this expression the other fields can be calculated using Eqn. 2.66 and Eqn. 2.70.

Lastly, using the result of Eqn. 2.66 we can make a statement about the effect of these beam-generated fields on the beam itself. Using the Lorentz force law (Eqn. 2.7) and calculating the transverse component by substituting $\boldsymbol{B}$ with Eqn. 2.66 gives

$$F_r = \frac{q}{\gamma_L^2}E_r. \tag{2.73}$$

This tells us that the forces of the focusing self-magnetic field and repulsive self-electric field cancel out to order $1/\gamma_L^2$. Thus, for ultra-relativistic beams we can safely ignore self-fields when modeling beam dynamics.

2.2 Plasma

Plasma is a unique state of matter characterized by extremely high temperatures. In this state, electrons are energetic enough to be free from binding to any one nucleus. As such, a plasma is made up of negatively charged electrons and positively charged ions where both populations are subject to external electromagnetic fields. Meanwhile, if the charge balance between the electron and ion populations is equivalent, then the plasma itself is neutral and produces no electromagnetic fields when left unperturbed in equilibrium. This quasi-neutral assumption is common in plasma physics, and we make this assumption here in the case of an unperturbed plasma. For this section, we will be using information from Refs. [36, 37, 38] to define some basic plasma parameters and conditions for a plasma to behave under ideal conditions.

The total range of parameters for all plasmas is vast, so here we will limit ourselves to a smaller parameter range relevant for plasma wakefield acceleration. Generally, a higher density allows for greater electromagnetic fields in a plasma which we can use to contain and accelerate an electron beam. Therefore, the typical range of densities for plasma wakefield accelerators is between 10^{15} to 10^{19} cm^{-3}, where the density of air is approximately 3×10^{19} cm^{-3}. The plasmas we consider are ionized in a laboratory setting, and so they have temperatures that range on the order of 1 to 100 eV. The volume of the plasma we are considering are anywhere from a meter-scale plasma wakefield accelerator stage to a ~ 100 μm thickness for plasma lenses, and the transverse sizes of these plasma are on the order of a few-hundred μm.

2.2.1 Plasma Parameters

Since the electrons and ions in a plasma evolve according to electromagnetic fields, we can derive parameters to represent timescales and length scales for these interactions. If we consider a plasma with number densities of n_e for electrons, n_i for ions, and n_n for neutral particles, we calculate the displacement force for light electrons while assuming the heavier ions remain stationary. If a slab of electrons all displace by a distance x relative to the ions, then the electric field from

Gauss' Law is

$$E_x = \frac{en_e}{\epsilon_0}x. \tag{2.74}$$

The force on an electron in this scenario is $F_x = -eE_x$, and this leads to an equation of motion of

$$x'' = -\frac{n_e e^2}{m_e \epsilon_0}x. \tag{2.75}$$

The solution is a simple harmonic oscillator, where the frequency

$$\omega_p = \sqrt{\frac{n_e e^2}{m_e \epsilon_0}} \tag{2.76}$$

is the electron plasma frequency. A similar expression can be derived for ion plasma frequency ω_{pi}, with $e^2 \to Z^2 e^2$ for ions with a charge state of Z. However, since the ion masses are much larger than the electron masses, the ion frequency is relatively small compared to the electron frequency. For this reason, the ions are typically considered stationary.

For plasma wakes excited by a relativistic electron beam or a high intensity laser pulse (both of which travel at approximately c), we can derive a length scale using the time period of plasma oscillations with the speed of light:

$$\lambda_p = 2\pi \frac{c}{\omega_p}. \tag{2.77}$$

This is known as the plasma wavelength, and along with the plasma wavenumber $k_p = \omega_p/c$ give approximate length scales to plasma wakes driven by electron beams or lasers. As an example, if we consider the density range 10^{15} to 10^{19} cm^{-3}, this gives us a range of plasma wavelengths between 168 and 1.68 μm.

The first criterion for a plasma is that the collective electromagnetic effects dominate over collisions. This is true if the period of plasma motion is shorter than the average collision time between electrons and neutral particles τ_n, or

$$\omega_p \tau_n > 1. \tag{2.78}$$

This average collision time is

$$\tau_n = \frac{1}{n_n \sigma_n}\sqrt{\frac{m_e}{2k_B T_e}}, \tag{2.79}$$

where $\sigma_n \approx 10^{-19}$ m^2 is the neutral atomic cross section, k_B is the Boltzmann constant, and T_e is the electron temperature for the plasma. This is easily satisfied for plasmas we are interested in with densities ranging from 10^{15} to 10^{19} cm^{-3} and temperatures from 1 to 100 eV.

The second criterion for a plasma is that it is large enough to provide shielding for external electric fields. If an electric field is applied to a plasma, the plasma electrons will move to cancel out the field and subsequently shield the interior of the plasma. The length over which this process occurs is called the Debye length λ_D, and is approximately

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{n_e e^2}}. \tag{2.80}$$

If the plasma exists in a system with a length scale L, then the Debye length must be much smaller than the system size:

$$\lambda_D \ll L. \tag{2.81}$$

For the range of plasma densities and temperatures we consider here, this Debye length ranges from 2.4 nm at 10^{19} cm^{-3} density and 1 eV temperature, up to 2.4 μm at 10^{15} cm^{-3} density and 100 eV temperature. For a system length on the order of the plasma wavelength λ_p, the Debye length is always much smaller.

The third criterion for a plasma is that there are enough electrons N_D within a Debye sphere to effectively perform the shielding. This requirement is satisfied when

$$N_D = \frac{4}{3}\pi n_e \lambda_D^3 \gg 1. \tag{2.82}$$

If we look at the upper range of parameters when $n_e = 10^{19}$ cm^{-3}, for 1 eV there is less than one particle in the Debye sphere while there are 544 at 100 eV. For this criterion to be satisfied, at the higher densities the plasma temperature needs to be larger.

2.3 Plasma Wakes

With the introduction of the electron beam and the plasma out of the way, we move towards combining the two concepts for a scenario where the relativistic electron beam propagates into a

plasma and induces a perturbation. The electron beam here is referred to as a "driver", and the density response to this driver is referred to as the "wake". The nature of this plasma response depends on if the ratio of the plasma density to the electron beam density. If the plasma is much more dense than the electron beam, then the electron beam produces a density perturbation of the plasma electrons and we can solve the plasma response analytically through the linearization of the plasma fluid equations. If the electron beam is more dense than the plasma, then the plasma will exhibit a nonlinear response that is more difficult to model but produces more advantageous electromagnetic fields, or "wakefields," for electron beam acceleration and focusing.

Here I will assume the plasma wake is caused by an ultra-relativistic electron beam propagating through a plasma. This choice of a particle beam driver is called "Plasma Wakefield Acceleration" (PWFA) in the accelerator community. However, much of this formalism can be applied to an intense laser pulse driver rather than an electron beam driver. When a laser driver is used, the accelerator community refers to the process as "Laser Wakefield Acceleration" (LWFA) [39]. While PWFA formalism is used here, LWFA is a rich field of Physics with its own array of benefits and challenges.

2.3.1 Linear Plasma Wakes

The first regime of PWFA to be studied is that of the linear regime due to its analytic tractability. The linear regime refers to the small amplitude limit of oscillations in the plasma density. Here, we treat the electron beam driver as a perturbative force that results in a linear plasma response. To understand better what this means, we first introduce how linear waves are modeled within a plasma. In this section, we follow Keinigs and Jones [40] using additional information from Refs. [41, 42, 43].

For a cold fluid plasma where the ions are assumed stationary and the density is initially infinitely uniform, the plasma evolves according to the continuity equation:

$$\frac{\partial n_e}{\partial t} + \nabla \cdot (n_e \boldsymbol{v}) = 0. \tag{2.83}$$

This equation preserves the total amount of charge in the system if there are no sources or sinks of electrons.[7] We can also write the Lorentz force law to calculate the force on a given volume of electrons:

$$m_e \frac{\partial(n_e \boldsymbol{v})}{\partial t} = -e n_e (\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}). \tag{2.84}$$

Additionally, the charge density $\rho = e(n_i - n_e)$ and current density $\boldsymbol{J} = -e n_e \boldsymbol{v}$ can be used with Maxwell's equations (Eqns. 2.52-2.55) to calculate the electromagnetic response to fluid motion.

For the linear regime, the electron density is assumed to have small amplitude offsets from its unperturbed value, or $n_e = n_0 + \delta n$. We also assume that the plasma currents induced by such a response are small such that $v \ll c$. This leads to the magnetic field being negligibly small compared to the electric field. If we expand the continuity equation to first-order in δn we get

$$\frac{\partial \delta n}{\partial t} + n_0 \nabla \cdot \boldsymbol{v} = 0. \tag{2.85}$$

If we take the electric field component of Eqn. 2.84,

$$m_e \frac{\partial \boldsymbol{v}}{\partial t} = -e\boldsymbol{E}, \tag{2.86}$$

and Gauss' law (Eqn. 2.52) for the electric field due to the perturbed electron density,

$$\nabla \cdot \boldsymbol{E} = -\frac{e}{\epsilon_0} \delta n, \tag{2.87}$$

we can then take the divergence of Eqn. 2.86 and use substitutions for $\nabla \cdot \boldsymbol{E}$ and $\nabla \cdot \boldsymbol{v}$ to get

$$\frac{\partial^2 \delta n}{\partial t^2} + \frac{n_0 e^2}{m_e \epsilon_0} \delta n = 0. \tag{2.88}$$

This is an electrostatic wave with an oscillation frequency equivalent to the plasma frequency.

Next, we can take an ultra-relativistic electron beam to serve as the external source of this plasma wave. An electron beam of density n_b adds an additional term to Gauss' law, and Eqn. 2.88 is modified to

$$\frac{\partial^2 \delta n}{\partial t^2} + \omega_p^2 \delta n = -\omega_p^2 n_b. \tag{2.89}$$

[7] More electrons could be generated through additional ionization processes, as can electrons be recombined with ions to act as a sink.

The electrons in this beam have a relativistic energy given by γ_L, which increases their effective mass to $\gamma_L m_e$. Because the beam electrons are effectively much heavier than the plasma electrons, we make a quasistatic approximation that the beam distribution is constant on the time scale of the plasma response. Using the co-moving coordinate ξ introduced earlier in Eqn. 2.4, we can rewrite the equation of motion to a frame traveling at speed v_b:

$$\frac{\partial^2 \delta n}{\partial \xi^2} + \frac{\omega_p^2}{v_b^2}\delta n = -\frac{\omega_p^2}{v_b^2} n_b. \tag{2.90}$$

This equation produces a driven harmonic oscillator at each transverse position x and y independently of one another. To solve for the density perturbation at a given transverse position, we use the Green's function for Eqn. 2.90 (Appendix A of Ref. [33]):

$$G(\xi, \xi') = -\frac{\omega_p}{v_b} \sin\left[\frac{\omega_p}{v_b}(\xi - \xi')\right] \Theta(\xi' - \xi), \tag{2.91}$$

where Θ is the Heavyside step function and its inclusion leads to the wake being undriven in front of the electron beam. We can then calculate the density perturbation as

$$\delta n(x, y, \xi) = -\frac{\omega_p}{v_b} \int_{-\inf}^{\xi} d\xi' n_b(x, y, \xi') \sin\left[\frac{\omega_p}{v_b}(\xi - \xi')\right]. \tag{2.92}$$

Assuming the electron beam is ultra-relativistic with $v_b \approx c$, the quantity $\omega_p / v_b = k_p$ is the plasma wavenumber.

Equation 2.92 has two qualitatively different solutions depending on the length of the drive beam compared to the plasma wavelength λ_p. If the drive beam is much longer than λ_p, then the integral is integrated over many full periods of the sine function and the result tends to zero. Intuitively, this is the result of plasma shielding as the plasma electrons are able to cancel out the electric field of the drive beam in the time scale of the drive beam interaction. If instead the drive beam is much shorter than λ_p, the beam drives a sinusoidal wave that persists in the plasma even after the electron beam has passed by. These two qualitative solutions are shown in Fig. 2.6 for a long and short drive beam, respectively.

With this density perturbation in the plasma, we can then go and calculate the electromagnetic fields occur within this plasma wake. To do so, we return to the inhomogeneous wave

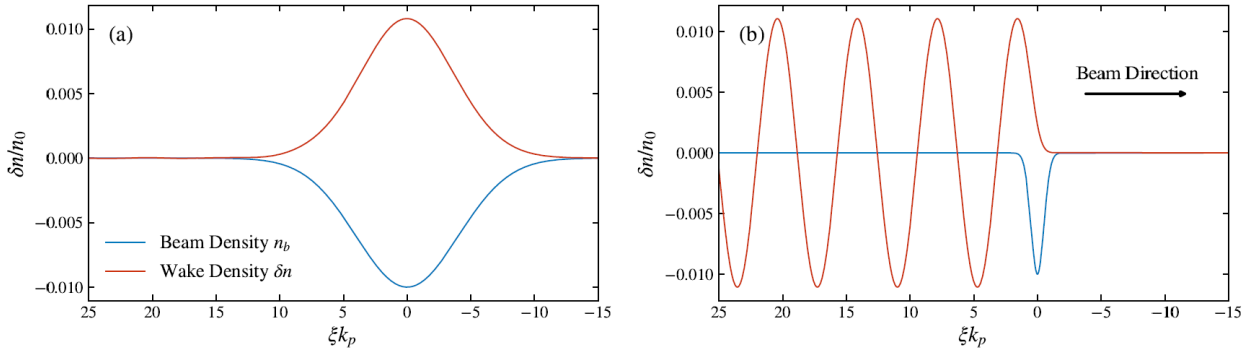

Figure 2.6: Linear plasma waves in density along the drive beam propagation axis, ξ. Subplot (a) shows the plasma density evolution when the drive beam is sufficiently long to allow for the plasma electrons to adiabatically screen the drive beam fields. Subplot (b) shows the plasma density wave driven by a short bunch. Here the plasma density wave continues even after the drive beam has passed. Reproduced from Ref. [43]

equations derived earlier for the electric field (Eqn. 2.59) and magnetic field (Eqn. 2.60). The source terms for the linear wake are $\rho = -e(\delta n + n_b)$ and $\boldsymbol{J} = -e(n_0\boldsymbol{v} + n_b c\hat{z})$, and the time derivative of $\boldsymbol{v}$ is given by the Lorentz force equation earlier in Eqn. 2.86. With these substitutions, the inhomogeneous wave equations become

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2} - k_p^2\right)\boldsymbol{E} = -\frac{e}{\epsilon_0}\nabla(\delta n + n_b) - \frac{e}{\epsilon_0 c^2}c\hat{z}\frac{\partial n_b}{\partial t} \tag{2.93}$$

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2} - k_p^2\right)\boldsymbol{B} = -\mu_0 ec\nabla \times (n_b\hat{z}). \tag{2.94}$$

Also, as we did in Sec. 2.1.6, when we introduce ξ and transform to the co-moving frame of Eqn. 2.4 we cancel out the 2$^{\text{nd}}$ order time derivative terms with the longitudinal component of ∇^2. These inhomogeneous wave equations are now written in the co-moving coordinate system as

$$(\nabla_\perp^2 - k_p^2)\boldsymbol{E} = -\frac{e}{\epsilon_0}\nabla\delta n - \frac{e}{\epsilon_0}\nabla_\perp n_b \tag{2.95}$$

$$(\nabla_\perp^2 - k_p^2)\boldsymbol{B} = -\mu_0 ec\left(\frac{\partial n_b}{\partial y}\hat{x} - \frac{\partial n_b}{\partial x}\hat{y}\right), \tag{2.96}$$

which is nearly the same as what we had in Eqn. 2.63 and Eqn. 2.64 but with additional source term in the electric field equation for the plasma density perturbation.

As we discovered in Sec. 2.1.6, the self-fields from the electron beam cancel out for ultra-relativistic beams with large γ_L^2. With this in mind, we can ignore the fields from the electron beam

and and focus on just the source term for the plasma wake.[8] This leaves us with

$$(\nabla_\perp^2 - k_p^2)\boldsymbol{E} = -\frac{e}{\epsilon_0}\nabla\delta n \tag{2.97}$$

$$(\nabla_\perp^2 - k_p^2)\boldsymbol{B} = 0. \tag{2.98}$$

The magnetic fields from this linear wake are simply zero, and the electric fields can be found by independently solving the longitudinal and transverse dimensions. For simplicity, we assume that the electron beam that drove the wake was cylindrically symmetric and the resulting wake is also cylindrically symmetric. The field is also assumed to be separable such that $E_{z,r} = R_{z,r}(r)Z_{z,r}(\xi)$, which implies that both the electron beam n_b and density perturbation δn can also be treated as this way.

The solution to Eqn. 2.97 requires the use of Green's function in cylindrical coordinates. The longitudinal electric field produced by the linear wake is

$$E_z(r,\xi) = \frac{e}{\epsilon_0}\int_0^\infty r'dr'\frac{\partial}{\partial\xi}\delta n(r',\xi)K_0(k_p r_>)I_0(k_p r_<) \tag{2.99}$$

and the transverse electric field is

$$E_r(r,\xi) = \frac{e}{\epsilon_0}\int_0^\infty r'dr'\frac{\partial}{\partial r'}\delta n(r',\xi)K_1(k_p r_>)I_1(k_p r_<), \tag{2.100}$$

where K_n and I_n are modified Bessel functions, $r_>$ is the greater of r and r', and $r_<$ is the lesser of the two.

These electric fields are shown in Fig. 2.7 for the case of a Gaussian drive beam (Eqn. 2.71) with transverse size $\sigma_x = \sigma_y = 1/k_p$, longitudinal size $\sigma_\xi = 0.5/k_p$, and a ratio of peak electron beam density to plasma density of $n_{b0}/n_0 = 0.01$. Qualitatively, we can see that the radial electric field is in phase with the density perturbation, while on the other hand the longitudinal field is $\pi/2$ out of phase. This arises from the different partial derivatives in Eqns. 2.99 and 2.100 acting on the sine function in δn of Eqn. 2.92. The exact shape of these fields becomes important once we consider the implications of adding a second particle beam within this plasma wake.

[8] The electron beam fields aren't zero, but since these equations are linear we can take advantage of the fact that the total fields are the sum from both sources. If needed, we can sum the beam fields from Sec. 2.1.6 with the fields from the wake here.

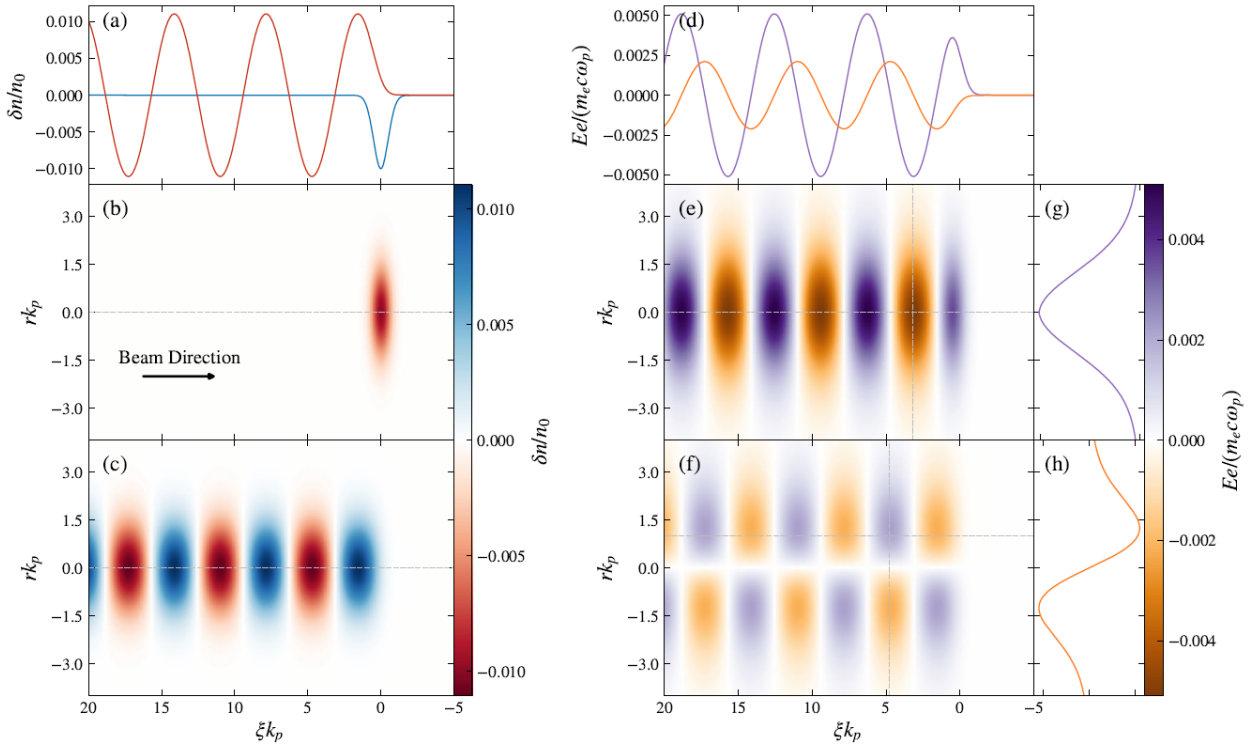

Figure 2.7: Plots of the density wave and resulting wakefields in the linear regime. (a) is the on-axis drive beam and plasma density, (b) is a view of the drive beam density in the transverse-longitudinal plane, (c) is a view of the linear plasma wave in the same plane behind the drive beam. (d) plots longitudinal lineouts of the on-axis longitudinal (purple) and off-axis transverse (orange) wakefields. The oscillations of the two are $\pi/2$ out of phase of one another. (e) and (g) plot the longitudinal wakefield behind the drive beam and its transverse variation. (f) and (h) plot the same for the transverse wakefield. Dashed lines denote where the lineouts are taken in their respective subplots. Reproduced from Ref. [33]

2.3.2 Linear Plasma Wakefield Accelerator

With the electric fields for a linear plasma wake derived in Sec. 2.3.1, we can now consider how we may use these fields. If we have a second particle beam placed behind the first, drive beam, then this second particle beam will "witness" the wakefields that arise in the plasma. This second particle beam is called the "witness beam", and it is this beam that we primarily consider when discussing the effects of a plasma wakefield device. We can either optimally place this beam in a phase of the generated longitudinal electric fields for maximum acceleration, as is the case with a plasma wakefield accelerator. Or we can consider using the radial electric field to provide a focusing

force to the witness beam in a short plasma, as will the case when we introduce passive plasma lenses in Chapter 3. Both of these processes are tied to the witness beam's interaction with the plasma wake, and so the discussion of plasma lenses heavily includes an examination of the plasma wakefield accelerator.

For an accelerator application, there are three things we must consider: (1) the accelerating gradient itself, (2) that the accelerating longitudinal field is uniform throughout the length of the witness beam to preserve the beam's longitudinal quality, and (3) the transverse beam dynamics of the witness beam preserve the beam's transverse quality. In this section we briefly go over these three aspects in the context of a linear plasma wakefield accelerator.

First, the maximum accelerating gradient can be obtained by simply choosing the drive beam and plasma parameters in such a way that the longitudinal field of Eqn. 2.99 is maximized. Most of the choice in parameter space here revolves around the drive beam, as the plasma needs to just be more dense to remain in the linear regime. In particular, the longitudinal current profile of the beam $Z_b(\xi)$ plays a large role in the shape of the plasma wake. If we follow Katsouleas et. al. [44], the longitudinal field can be written

$$E_z(r,\xi) = -\frac{e}{\epsilon_0} R(r) \int_{\infty}^{\xi} d\xi' Z_b(\xi') \cos[k_p(\xi - \xi')], \tag{2.101}$$

where $R(r)$ is the radial component of Eqn. 2.99. Then we consider the simple case of a cylindrical beam with uniform density. For such a beam with charge Q, length $\Delta\xi$, radius Δr, and density $n_{b0} = \frac{Q}{e\pi\Delta\xi\Delta r^2}$, the beam current profile is $Z_b(\xi) = n_{b0}$ within the bounds of the drive beam and zero otherwise. Within the beam, the longitudinal wakefield is

$$E_z(r,\xi) = -\frac{e}{\epsilon_0} R(r) \frac{1}{k_p} n_{b0} \sin(k_p \xi). \tag{2.102}$$

To efficiently drive a wake in the plasma, the length of the beam should be smaller than a quarter-phase of the sine function, or $\Delta\xi < \lambda_p/2$. Behind the drive beam, the longitudinal wakefield is

$$E_z(r,\xi) = -\frac{e}{\epsilon_0} R(r) \frac{1}{k_p} n_{b0} \sqrt{2 - 2\cos(k_p \Delta\xi)} \cos(k_p \xi + \phi), \tag{2.103}$$

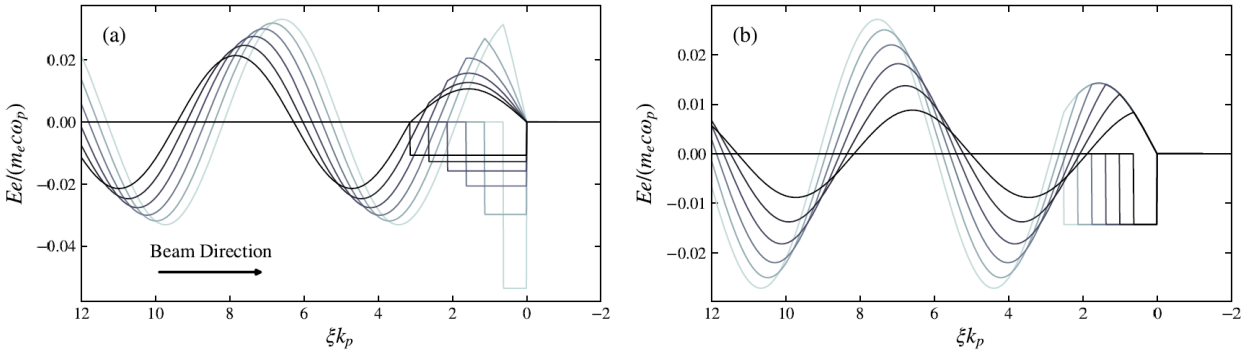

Figure 2.8: Variation of linear regime longitudinal wakefields for various current profiles. The rectangles represent the cylindrical drive beam current profiles. (a) shows the beam charge held constant while the current is varied through the bunch length. Higher current produces a stronger wake. (b) shows the beam current held constant while the total charge is varied through the bunch length. More charge also produces a stronger wake. Reproduced from Ref. [33]

where ϕ is the wake's phase at the tail end of the drive beam:

$$\phi = \arctan\left[-\frac{1-\cos(k_p\Delta\xi)}{\sin(k_p\Delta\xi)}\right]. \tag{2.104}$$

The longitudinal wakefield is largest at $r = 0$, so the wakefield amplitude is

$$E_{z0} = -\frac{e}{\epsilon_0}R(0)\frac{1}{k_p}n_{b0}\sqrt{2-2\cos(k_p\Delta\xi)}. \tag{2.105}$$

The longitudinal wakefield produced by such a cylindrical drive beam is plotted in Fig. 2.8 for different values of charge and current. In general, to produce the strongest accelerating fields you would want a drive beam with both high charge and a small length for high current.

The second consideration, a constant accelerating field across a witness beam, is more difficult to approach if your witness beam has a finite length. As we have seen in Figs. 2.7 and 2.8, the longitudinal wakefield is a sinusoidal function that is most certainly not constant. However, just as the plasma wake arises from the drive beam inducing a plasma response, the witness beam can also subsequently invoke a plasma response in the wake that "flattens" the longitudinal wakefield profile to give a constant, accelerating field across its duration. This is known as "beam loading," and proper beam loading requires the current profile of the witness beam to be carefully tailored to the drive beam's wake.

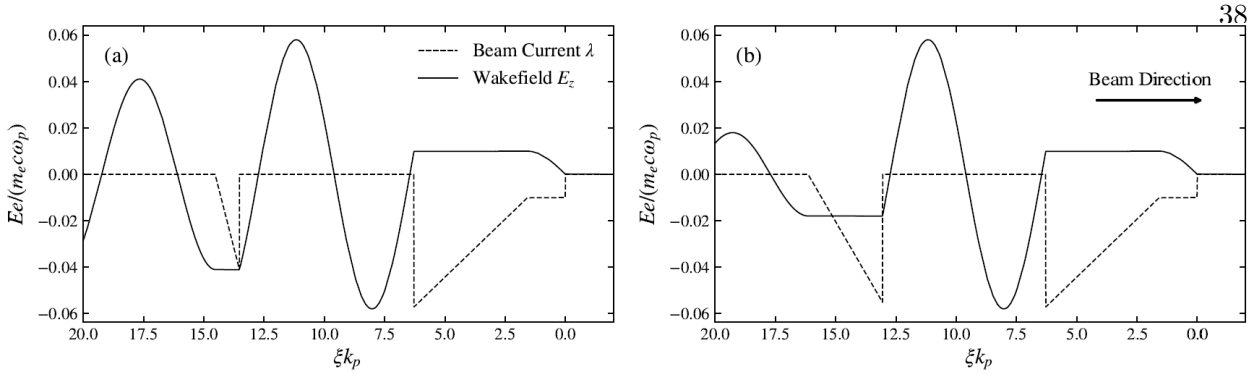

Figure 2.9: The longitudinal wakefield (solid) and beam current (dashed) for an optimally loaded plasma wakefield accelerator. The beam current plots include both the drive beam (smaller ξ) and witness beam (larger ξ). The drive beam has a trapezoidal current profile with an initial step in order to produce a constant deceleration field within the drive beam, while the witness beam has a triangular current profile to flatten the longitudinal wakefield for constant acceleration. (a) is an example of a witness beam with low charge, and this beam can be placed in a region of high acceleration. (b) shows a witness beam with more charge, and this beam needs to be located in a region of less acceleration than in (a). Reproduced from Ref. [33]

The wake produced by the combined current of a drive and witness beam can be modelled in the linear regime by including both beams in the beam density using $n_b(r,\xi) = n_d(r,\xi) + n_w(r,\xi)$. Then, the wake is calculated as before in Sec. 2.3.1 with both beams contributing to the wake. In this linear regime the solution for optimal beam loading is to shape the beams to have a trapezoidal[9] current profile, as seen in Fig. 2.9. By flatten the wake for optimal beam loading, the price is a lower accelerating gradient depending on how much charge is in the witness beam. More charge requires more beam loading which leads to a smaller accelerating gradient overall.

The third consideration for the linear plasma wakefield accelerator has to do with the radial components of both the longitudinal accelerating field and the transverse focusing field, and this is where the linear regime fails to impress. The radial component of Eqn. 2.99 means that a witness beam with a finite radial size would experience different accelerating fields across its width. The witness beam would then need to be a narrow as possible to minimize this spread, but even then it is still difficult to achieve optimal beam loading for all r.[10] The radial dependence of the transverse fields of Eqn. 2.100 are also not ideal for electron beams since they are nonlinear in r. Nonlinear

[9] Or triangular if the current goes to zero on the tail edge.

[10] Not to mention that trapezoidal beams are much harder to create than Gaussian beams.

focusing forces introduce aberrations into the beam that reduce the quality, and for the beam dynamics covered in Sec. 2.1.2 we would want a focusing force that is linear in r for Hill's equation to be applicable. Luckily for us, there is another regime beyond the linear regime where these concerns are alleviated.

2.3.3 Nonlinear Plasma Wakes

The condition for the linear plasma wake regime was for the drive beam density to be less than the plasma density. This condition resulted in the plasma responding to the drive beam perturbation in a linear response. As we increase the drive beam density further and further, this assumption begins to break down. When the drive beam density is larger than the plasma density,[11] the plasma electrons respond to this intense driving force by being completely expelled from the wake's center and leaving behind a region of stationary ions. It is this uniform density of stationary ions that provides some beneficial qualities to the wakefields in this nonlinear regime.

In the linear regime, the longitudinal wakefield had an amplitude of $E_{z0} \approx \frac{e\delta n}{\epsilon_0 k_p} = \frac{\delta_n}{n_0}\frac{m_e c\omega_p}{e}$ where the charge separation is approximately the density of the drive beam $\delta_n = n_{b0}$. If we take this expression to the absolute limit of complete charge separation where $\delta_n = n_0$, we get a maximum electric field:

$$E_z = \frac{m_e c\omega_p}{e} = \sqrt{\frac{m_e c^2 n_0}{\epsilon_0}}. \tag{2.106}$$

This limit is called the wave breaking field [45], and in the nonlinear regime we can expect longitudinal accelerating fields on the order of this wave breaking field. In the range of densities we are considering for PWFA's, a density range of 10^{16} to 10^{18} cm^{-3} results in a range of accelerating fields from 10 to 100 GV/m. These fields are much larger than those of the linear regime, and also orders of magnitude larger than conventional acceleration techniques.

Repeating the methods used in the derivation of the plasma wakes in the linear regime, first the plasma electron response is modeled to get the wake density profile and then from this profile

[11] The exact condition is vague. Some sources cite $n_b > n_0/2$, some cite $n_b > n_0$, and some cite $n_b > 2n_0$. Roughly between the linear and nonlinear regimes in this blurry region is a "quasi-nonlinear" regime that exhibits characteristic of both.

we can get the wakefields. However, unlike the linear regime this nonlinear regime is more difficult to model. Here we present the results of Lu et. al., [46, 47] and of [48] for the analytic effort in modeling this regime, and highlight the relevant results.

From Ref. [47], the nonlinear regime is reached when the trajectories of plasma electrons cross one another when the dense drive beam passes through. Mathematically, this would imply that there exists at least two plasma electrons of initial radii r_1 and r_2 and at $\xi < 0$ they are arranged $r_1 > r_2$. Then, after the electron beam passes through for $\xi > 0$, the innermost plasma electron was blown out further so that $r_1 < r_2$. It is then possible to work out the maximum blowout radius for each electron r_m as a function of its initial distance from the drive beam axis r_0. Since the length scale of plasma motion is dictated by the plasma frequency, the location of this maximum blowout radius will be at approximately the same longitudinal value ξ for all plasma electrons. Dawson in Ref. [49] found an expression to model this blowout radius in the non-relativistic limit by treating the plasma electrons as a series of concentric electron rings:

$$\frac{d^2r}{d\xi^2} = -\frac{k_p r}{2} + k_p c(r_0, r, \xi r, \tag{2.107}$$

where $-k_p r/2$ is the force due to the stationary, uniform ion column and

$$c(r_0, r, \xi) = \frac{1}{2}r_0^2 + \frac{1}{n_0}\int_0^r r' dr' n_b(r' x \xi) \tag{2.108}$$

is the force from the electron beam if there was no trajectory crossing. The threshold for trajectory crossing can then be solved by finding beam parameters which result in $\frac{dr_m}{dr_0} = 0$. Figure 2.10 shows electron trajectories for the case of a uniform, cylindrical drive beam of radius a[47]. Trajectory crossing is reached at $n_b > 1.792 n_0$ for this beam distribution, where this crossing first appears for electrons with initial radius $r_0 = a$. As the drive beam density increases, the the threshold for trajectory crossing lowers until it is just about at the drive beam axis, at which point all of the plasma electrons are blown out of the wake. This is the ideal state of a wake in the nonlinear regime, and why it is also common to refer to this regime as the "blowout regime."

The plasma density profile of the blowout regime is investigate further by using particle in cell (PIC) simulations to numerically solve the plasma evolution. PIC itself is introduced further in

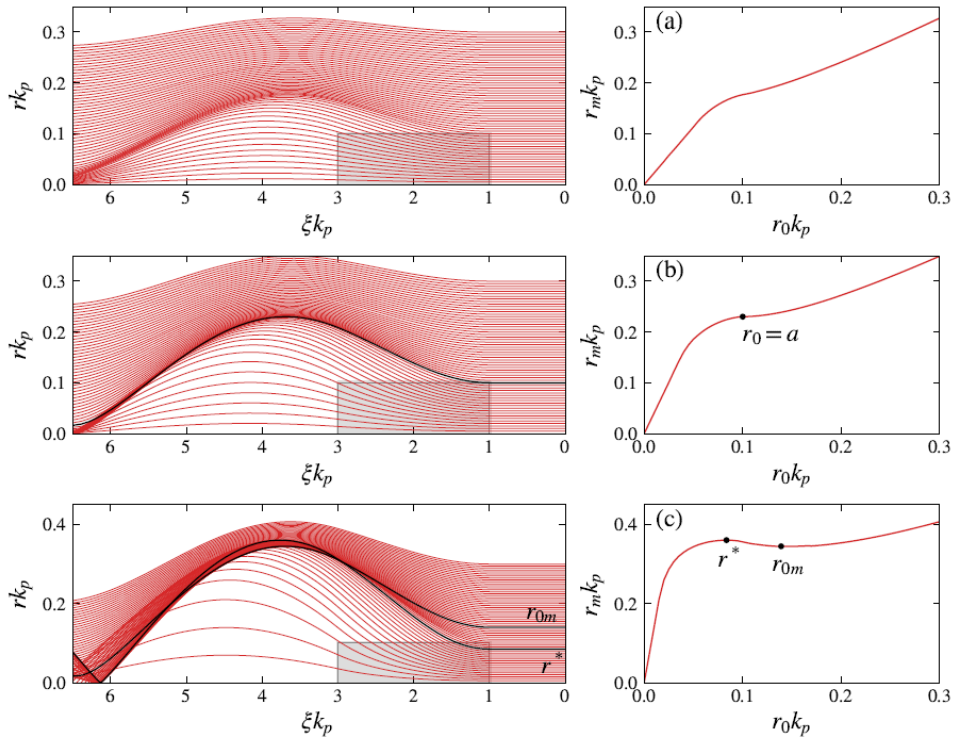

Figure 2.10: Plasma electron trajectories for different radial start positions. The shaded box represents the electron drive beam, and the right subplots show the maximum blowout radius vs start position for each electron. Trajectory crossing occurs if the slope of this relation reaches zero or below. (a) shows trajectories when $n_b = n_0$ and no trajectory crossing occurs. (b) shows $n_b = 1.792 n_0$ and the threshold is reached for an electron with $r_0 = a$. (c) shows $n_b = 4 n_0$ where trajectory crossing occurs for a range of plasma electron start positions. Reproduced from Ref. [47]

Sec. 2.4.1. In these numerical studies, the blowout wake is observed to have three separate regions. First there is the blowout itself, characterized by the stationary ion column and lack of any plasma electrons. Second is the high density electron sheath that surrounds this blowout, comprised of the plasma electrons which were expelled from the wake. Third is the region in the plasma beyond the electron sheath, where the plasma is far enough away from the drive beam to linearly respond to the drive beam and electron sheath. These regions are visible in Fig. 2.11, and the boundary between the blowout and the sheath is a sharp, discontinuous transition.

The sharp sheath boundary can be described by a function $r_b(\xi)$ of the longitudinal coordinate ξ, and this radius is approximately that of the electron with the smallest deflection from the drive beam axis. In Ref. [46, 47], the single particle equation of motion can be solved for this innermost

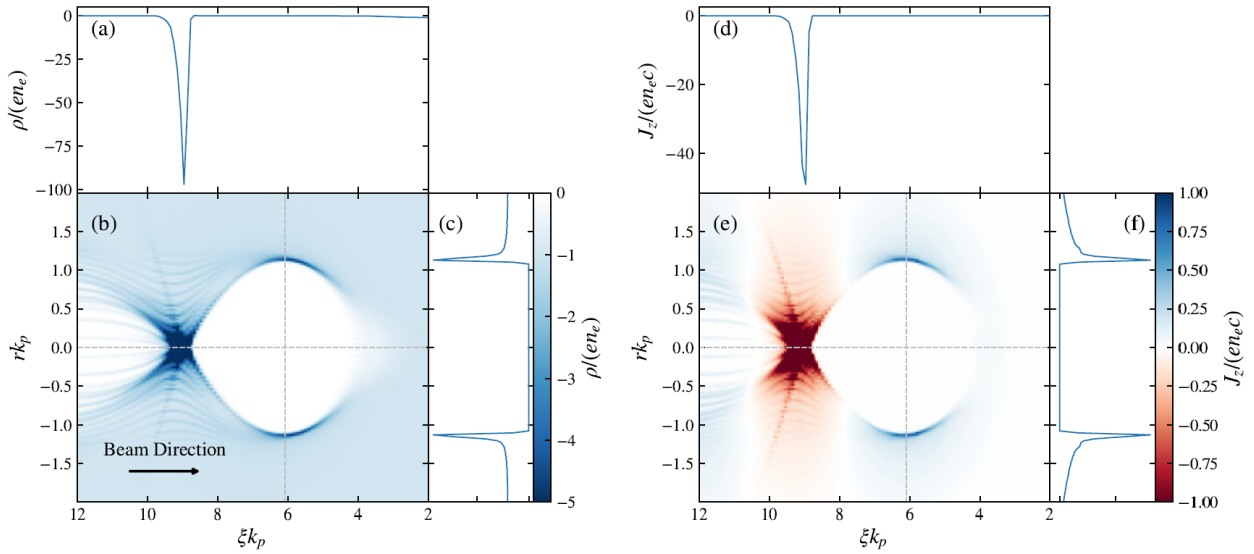

Figure 2.11: Structure of the plasma electron sheath density (b) and current (e) in the nonlinear blowout regime. (a) and (c) are lineouts of the plasma density across the center of the wake, whereas (d) and (f) are lineouts of the plasma's longitudinal current. For both, the plasma density and current are concentrated in the narrow sheath around the blowout. Reproduced from Ref. [33]

electron with initial radius r_{0m} to find the approximate blowout radius. As part of this derivation, a narrow profile for the sheath needs to be assumed. The requirement on this profile is that is satisfies the continuity equation

$$\frac{d}{d\xi}\int r dr\left(\rho - \frac{J_z}{c}\right) = 0. \tag{2.109}$$

Then, using the electromagnetic fields we'll discuss in Sec. 2.3.4, the inner radius of the blowout is described by

$$A(r_b)\frac{d^2 r_b}{d\xi^2} + k_p B(r_b) r_b \left(\frac{dr_b}{d\xi}\right)^2 + k_p C(r_b) r_b = -\frac{k_p}{2\pi c e n_0}\frac{\lambda(\xi, r_b)}{r_b}, \tag{2.110}$$

with the beam current as the source term:

$$\lambda(\xi, r_b) = 2\pi c \int_0^r \rho_b(\xi, r') r' dr'. \tag{2.111}$$

The beam charge density ρ_b includes both the drive and witness beam in a two-bunch configuration. The functions $A(r_b)$, $B(r_b)$, and $C(r_b)$ depend on the assumptions made on the sheath profile [50, 51] and on the radial density profile of the plasma [52, 53]. For the case of a transversely uniform

plasma, these are

$$A(r_b) = 1 + k_p^2 \frac{r_b^2}{4} + k_p^2 \frac{r_b^2 \beta}{2} + k_p^2 \frac{r_b^3}{8} \frac{d\beta}{dr_n}, \tag{2.112}$$

$$B(r_b) = \frac{1}{2} + \frac{3\beta}{4} + \frac{3r_b}{4} \frac{d\beta}{dr_b} + \frac{r_b^2}{8} d^2\beta dr_b^2, \tag{2.113}$$

$$C(r_b) = \frac{1 + (1 + r_b^2 k_p^2 \beta/4)^2}{4(1 + r_b^2 k_p^2 \beta/4)^2}. \tag{2.114}$$

The function $\beta(r_b)$[12] depends on the sheath profile used, and so furthering the effort of making this model as simple as possible we use a sheath with a uniform density and width Δ. The continuity equation (Eqn. 2.109) determines the exact density:

$$\rho - \frac{J_z}{c} = -en_0 \frac{r_b^2}{(r_b + \Delta)^2 - r_b^2}. \tag{2.115}$$

Next, $\beta(r_b)$ can be calculated from the wake potential. Without going into too much detail here, the result for this function is

$$\beta(r_b) = \frac{(1+\alpha)^2 \ln(1+\alpha)^2}{(1-\alpha)^2 - 1} - 1; \quad \alpha = \frac{\Delta}{r_b}. \tag{2.116}$$

Finally, we have a full description for the blowout radius of Eqn. 2.110 for the simplest case.

This solution is difficult to use, but we can trim off some of the finer details by making some valid assumptions. First, for an accelerator application in we would like for the wake to be strongly driven. We can assume the ultra-relativistic limit for the maximum blowout radius such that $r_m \gg k_p^{-1}$. Next, we can assume that the sheath width is much smaller than the blowout radius, $\Delta \ll r_b$, which is a safe assumption to make considering the sharp sheaths observed in simulations. This assumption lets us approximate $\beta(r_b) \approx 2\alpha(r_b)$. Dropping more smaller quantities gives the following:

$$A(r_b) = k_p^2 \frac{r_b^2}{4}, \quad B(r_b) = \frac{1}{2}, \quad C(r_b) = \frac{1}{4}. \tag{2.117}$$

Using these simpler functions, our differential equation for the sheath blowout radius becomes

$$k_p r_b \frac{d^2 r_b}{d\xi^2} + 2\left(\frac{dr_b}{d\xi}\right)^2 + 1 = -\frac{1}{2\pi c e n_0} \frac{4\lambda(\xi, r_b)}{r_b^2}. \tag{2.118}$$

[12] Sorry, this β is neither the CS parameter nor the ratio of v_b/c. We won't be needing this β for long, though.

Importantly, sheath width Δ does not appear in Eqn. 2.118. This means that dynamics of the sheath's blowout radius are unaffected by the details of the sheath profile.

Equation 2.118 is valid for the bulk of the wake, but at the very front and back of the blowout the sheath radius goes to zero and our assumption of $r_0 \gg k_p^{-1}$ falls apart. Fortunately, this result can still be salvaged. Using an energy-based argument, Golovanov et. al. showed in Ref. [48] that Eqn. 2.118 is still valid at the front of the wake, but since $\frac{d^2 r_b}{d\xi^2}$ diverges at the origin the equation is still numerically intractable. But, by making the substitution

$$\psi_\xi(\xi) = \frac{r_b^2(\xi)}{4}; \qquad \psi_\xi(\xi_0) = 0, \qquad \frac{1}{\psi_\xi}\left(\frac{d\psi_\xi}{d\xi}\right)^2\bigg|_{\xi=\xi_0} = 0. \tag{2.119}$$

we can write a form for the blowout radius differential equation that can be numerically integrated:

$$\frac{d^2\psi_\xi}{d\xi^2} + \left(\frac{2}{k_p} - 1\right)\frac{1}{2\psi_\xi}\left(\frac{d\psi_\xi}{d\xi}\right)^2 = -\frac{1}{2\pi c e n_0}\frac{\lambda(\xi, 2\sqrt{\psi_\xi}}{2k_p\psi_\xi} - \frac{1}{2k_p}. \tag{2.120}$$

Even here, special attention must be paid towards the head of the drive beam where both the current and blowout radius is zero. To account for this ξ_0 is chosen so that $\rho(\xi_0, r) = 1/2$. Comparisons between Eqn. 2.120 and PIC simulation results are shown in Fig. 2.12, and it is noticeable that the agreement between simulation and theory is better as the beam drives a more and more relativistic wake.

2.3.4 Nonlinear Plasma Wakefields

Previously in Sec. 2.3.3 we derived the blowout radius evolution for a nonlinear plasma wake. Here, using this blowout radius we can solve for the electromagnetic fields within this plasma blowout. The longitudinal wakefield is important for determining the maximum accelerating gradient of a nonlinear plasma wakefield accelerator, while the transverse wakefield is is important for determining the transverse beam dynamics of a witness beam in such a wake. For a passive plasma lens, this transverse wakefield will be the focusing force of our electron beam optic.

We begin with the inhomogeneous wave equations in the Lorentz gauge, using the scalar and

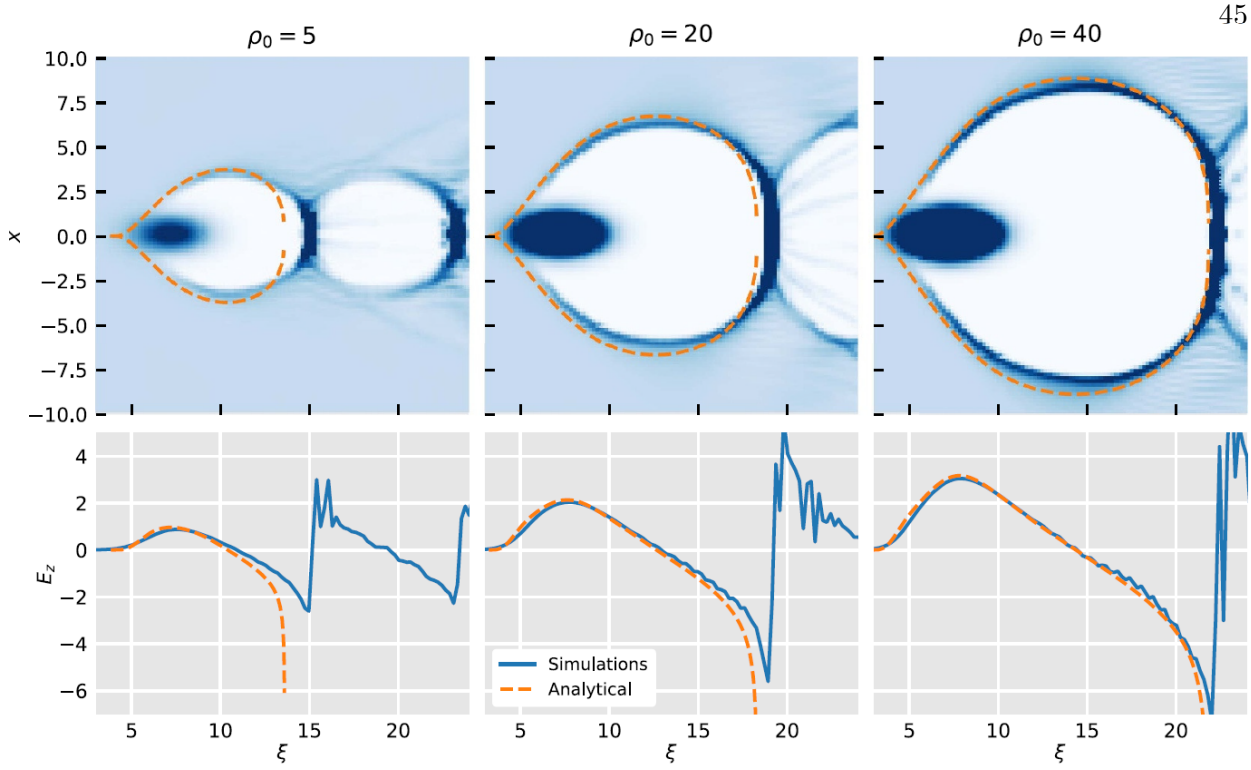

Figure 2.12: Theoretical predictions for the blowout radius (orange dashed) compared with simulation. The top plots show the plasma density profile where the blowout wake is clearly visible, and the drive beam increases in density from left to right and subsequently drives a stronger wake. The bottom plots are longitudinal lineouts of the longitudinal wakefield and also compare theory to simulation. Reproduced from Ref. [48]

vector potentials, φ and $\boldsymbol{A}$, respectively:

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)\varphi = -\frac{\rho}{\epsilon_0} \tag{2.121}$$

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\right)\boldsymbol{A} = -\mu_0 \boldsymbol{J}. \tag{2.122}$$

In the Lorentz gauge, by definition these potentials follow the relation

$$\nabla \cdot \boldsymbol{A} + \frac{1}{c^2}\frac{\partial \varphi}{\partial t} = 0. \tag{2.123}$$

As before, we want to take these three equations and transform them into the co-moving frame with the electron beam using longitudinal coordinate ξ (Eqn. 2.4). Doing so results in the longitudinal components of ∇^2 and $\partial^2/\partial t^2$ in these three equations cancelling out and we are left with the

following for the wave equations and Lorentz gauge relation:

$$\nabla_\perp^2 \varphi = -\frac{\rho}{\epsilon_0} \tag{2.124}$$

$$\nabla_\perp^2 \boldsymbol{A} = -\mu_0 \boldsymbol{J} \tag{2.125}$$

$$\nabla_\perp \cdot \boldsymbol{A}_\perp = -\frac{1}{c}\frac{\partial}{\partial \xi}(\varphi - cA_z). \tag{2.126}$$

Here we will define a wake potential,

$$\psi \equiv \varphi - cA_z, \tag{2.127}$$

and substitute this potential in after subtracting the longitudinal component of Eqn. 2.125 from Eqn. 2.124. This gives us a wave equation for the wake potential:

$$\nabla_\perp^2 \psi = \frac{1}{\epsilon_0}\left(\rho - \frac{J_z}{c}\right). \tag{2.128}$$

Now, our general approach is using the density and current profile of the plasma wake, $(\rho - J_z/c)$, to solve for the wakefield potential in Eqn. 2.128. Then, the electromagnetic fields can be found using the definitions of the vector and scalar potentials:

$$\boldsymbol{E} = -\nabla\varphi - c\frac{\partial \boldsymbol{A}}{\partial t} \tag{2.129}$$

$$\boldsymbol{B} = \nabla \times \boldsymbol{A}. \tag{2.130}$$

If we are solving for the longitudinal wakefield E_z, Eqn. 2.129 simplifies even further in the co-moving frame:

$$E_z = \frac{\partial \psi}{\partial \xi}. \tag{2.131}$$

Using the description of the blowout wake we found in Sec. 2.3.3, the interior of the wake is only populated by stationary ions so the density and current are simply $\rho = en_0$ and $J_z = 0$. If the blowout wake is cylindrically symmetric, the potentials within the blowout wake for $r < r_b\xi$ can be found as

$$\varphi = \varphi_0(\xi) - \frac{r^2}{4}\frac{en_0}{\epsilon_0} \tag{2.132}$$

$$A_z = A_{z0}(\xi) \tag{2.133}$$

$$A_r = -r\frac{1}{2c}\frac{d\psi_0}{d\xi} \tag{2.134}$$

$$\psi = \psi_0(\xi) - \frac{r^2}{4}\frac{en_0}{\epsilon_0} \tag{2.135}$$

where $\varphi_0(\xi)$, $A_{z0}(\xi)$, and $\psi_0(\xi)$ are their respective potential values on the $r = 0$ axis. We can assume that far from the blowout the potential is zero, or $\psi(\xi \to \infty) = 0$. The on-axis wakefield potential is then found by integrating Eqn. 2.128:

$$\psi_0 = \int_0^\infty \frac{dr}{r}\int_0^r r'dr'\frac{1}{\epsilon_0}\left(\rho - \frac{J_z}{c}\right). \tag{2.136}$$

To solve this integral we need to assume the electron sheath profile. Using the uniform density model from Eqn. 2.115, the integral is evaluated to give

$$\psi_0(\xi) = \frac{en_0}{\epsilon_0}\frac{r_b^2(\xi)}{4}(1 + \beta(\xi)), \tag{2.137}$$

where $\beta(\xi)$ is defined in Eqn. 2.116. In the ultra-relativistic approximation where $\beta \approx \alpha$ and $\alpha \ll 1$, the wake potential is

$$\psi(\xi, r) = \frac{en_0}{\epsilon_0}\left(\frac{r_b^2(\xi)}{4} - \frac{r^2}{4}\right). \tag{2.138}$$

The longitudinal wakefield can then be solved from Eqn. 2.131 as

$$E_z = \frac{en_0}{\epsilon_0}\frac{r_b}{2}\frac{dr_b}{d\xi}, \tag{2.139}$$

where we can see that E_z only depends on $r_b(\xi)$ and the plasma density n_0. In this nonlinear regime, E_z is nearly a linear function with a slope of $en_0/2\epsilon_0$, and the largest amplitude is close to the rear of the wake (Fig. 2.12). If we are trying to optimize a PWFA in the nonlinear regime, we would want to place a witness beam as close as possible to the rear of the wake, while also designing the witness beam's charge current and placement to flatten E_z with beam loading for uniform acceleration across the witness beam. Revisit Sec. 2.3.2 for more on beam loading. Flattening the longitudinal wakefield works much the same in the linear and nonlinear regimes, and the nonlinear wake is optimally loaded by a witness beam with a "trapezoidal" current profile [54].

The transverse wakefield is given by Lorentz force law,

$$\boldsymbol{W}_\perp = \boldsymbol{E}_\perp + c\hat{z} \times \boldsymbol{B}_\perp, \tag{2.140}$$

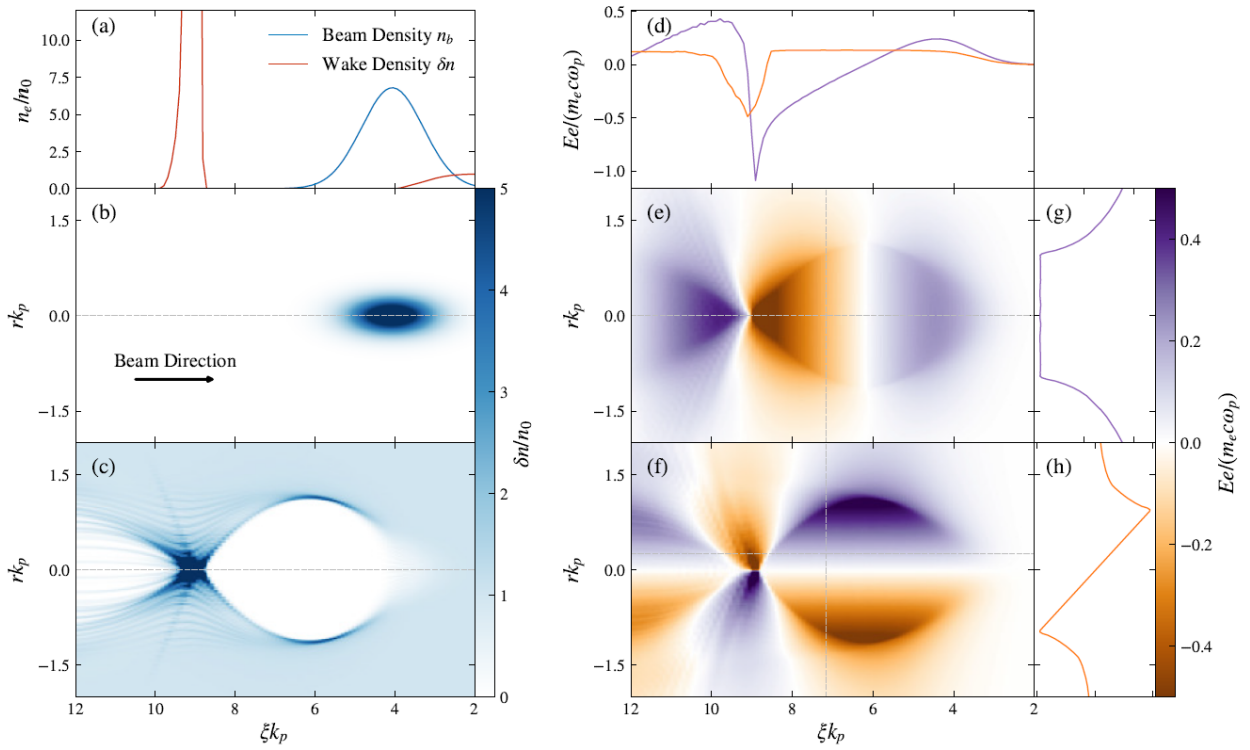

Figure 2.13: Plots of the plasma density response and resulting wakefields in the nonlinear blowout regime. (a) is the on-axis drive beam and plasma density, (b) is a view of the drive beam density in the transverse-longitudinal place, (c) is a view of the plasma density of the blowout wake. (d) plots longitudinal lineouts of the on-axis longitudinal (purple) and off-axis transverse (orange) wakefields. (e) and (g) plot the longitudinal wakefield behind the drive beam and its transverse variation. (f) and (h) plot the same for the transverse wakefield. Dashed lines denote where the lineouts are taken in their respective subplots. Reproduced from Ref. [33]

where we use $\boldsymbol{W}_\perp$ to denote the total transverse wakefield and can write the force on an electron as $\boldsymbol{F}_\perp = -e\boldsymbol{W}_\perp$. The transverse electric field is given by the transverse components of Eqn. 2.129 using the potentials within the blowout wake:

$$\boldsymbol{E}_\perp = \frac{1}{2} r \left(\frac{e n_0}{\epsilon_0} + \frac{d\psi_0}{d\xi} \right) \hat{r}. \tag{2.141}$$

The left term in the parenthesis is due to the uniform ion column, and the right term is the contribution from the radial plasma current. The magnetic field is also found using Eqn. 2.130 and the blowout potentials:

$$\boldsymbol{B}_\perp = \frac{1}{2} r \frac{1}{c} \frac{d\psi_0}{d\xi} \hat{\phi}, \tag{2.142}$$

which is entirely from the radial plasma current. When combining these two fields together, the radial plasma current from $\boldsymbol{E}_\perp$ and $\boldsymbol{B}_\perp$ cancel each other out. The transverse, focusing wakefield is then entirely due to the uniform ion column:

$$\boldsymbol{W}_\perp = E_{r,ion} = \frac{1}{2}\frac{en_0}{\epsilon}r. \tag{2.143}$$

We can make a few important observations here. Not only does Eqn. 2.143 only depend on the plasma density n_0, but if the plasma density is constant then this focusing force is a purely linear function in r. An ideal, linear focusing force results in no aberrations when focusing an electron beam. Also, this force is applicable to the Hill's equation (Eqn. 2.11) so the formalism we derived in Sec. 2.1 for an electron beam applies to beams that are subjected to this transverse force. Lastly, unlike a focusing quadrupole magnet, this ion column's electric field focuses simultaneously in both the x and y dimensions for an electron beam.

The wakefields in the nonlinear blowout regime are plotted in Fig. 2.13 for a typical example case. In this figure we can see another beneficial property of the nonlinear regime. The transverse profile of the longitudinal accelerating field is constant over the width of the blowout (Fig. 2.13(g)), which means that a witness beam will have all electrons accelerated at the same rate across its transverse phase space.

For the remainder of this dissertation we will be mostly concerned with the transverse wakefield. A thin, passive plasma lens is essentially a nonlinear blowout wake with a small longitudinal footprint, and so the the change in energy due to the longitudinal wakefields is negligible. We will introduce passive plasma lenses in greater detail in Chapter 3.

2.3.5 Emittance Preservation in Nonlinear PWFA

The last topic we will introduce for the nonlinear regime is that of emittance conservation in a PWFA. In Sec. 2.1.5 we quantified the transverse quality of an electron beam using the normalized emittance ϵ_N of Eqn. 2.49. This quantity is conserved for an electron beam undergoing both acceleration and linear focusing forces. At first glance we would look at the linear transverse wakefield of

Eqn. 2.143 and think that this shouldn't provide any issue for emittance preservation. However, we must include one key parameter when considering the focusing force on a relativistic electron: the longitudinal momentum. An electron beam will have some distribution of its longitudinal energy, and electrons with higher energy will effectively be more difficult to focus transversely. And so in a nonlinear PWFA with a witness beam with a large enough energy spread in a dense plasma with a strong focusing force, the emittance will naturally grow [55, 56].

This chromatic emittance growth arises from the inclusion of relativistic momentum when calculating the Lorentz force. This results in a Hill's equation that depends on both the plasma density and the relativistic Lorentz factor:

$$x'' + \frac{1}{2}\frac{e^2 n_0}{\epsilon_0}\frac{1}{\gamma_L m_e c^2}x = 0. \tag{2.144}$$

The equation for y is the same. In an electron beam with an energy spread in δ, we can solve for the transverse dynamics for each individual energy slice projection (Eqn. 2.34) using this Hill's equation for a nonlinear blowout. What we find is that the phase space ellipses of different energy projections rotate at different frequencies. After propagating in the plasma for some time, the total transverse projection of the beam will show the different energy slices spread out. As a result, the beam's emittance grows until it reaches a saturation value when the energy slices are distributed along all 2π of the phase in this betafunction oscillation. This is shown in Fig. 2.14 and demonstrates the key principle of this chromatic phase spreading: while the emittance values of each individual energy slice remain constant the emittance value of the full beam increases.

The solution to avoiding this emittance growth is to have an electron beam that keeps a constant size throughout this betafunction rotation. Qualitatively, if our initial beam's total projection in phase space is already a circle, then even if the different energy slices rotate at different frequencies then the emittance growth will be unnoticeable. This is achieved when our differential equation for the beam's betafunction, Eqn. 2.23, is solved for when $\beta'' = \beta' = 0$. If we assume that the plasma density is uniform, the solution is straightforward and reveals

$$\beta_m = \sqrt{\frac{1}{K}} = \sqrt{\frac{2\epsilon_0}{e^2 n_0}\gamma_L m c^2} = \frac{c}{\omega_p}\sqrt{2\gamma_L}, \tag{2.145}$$

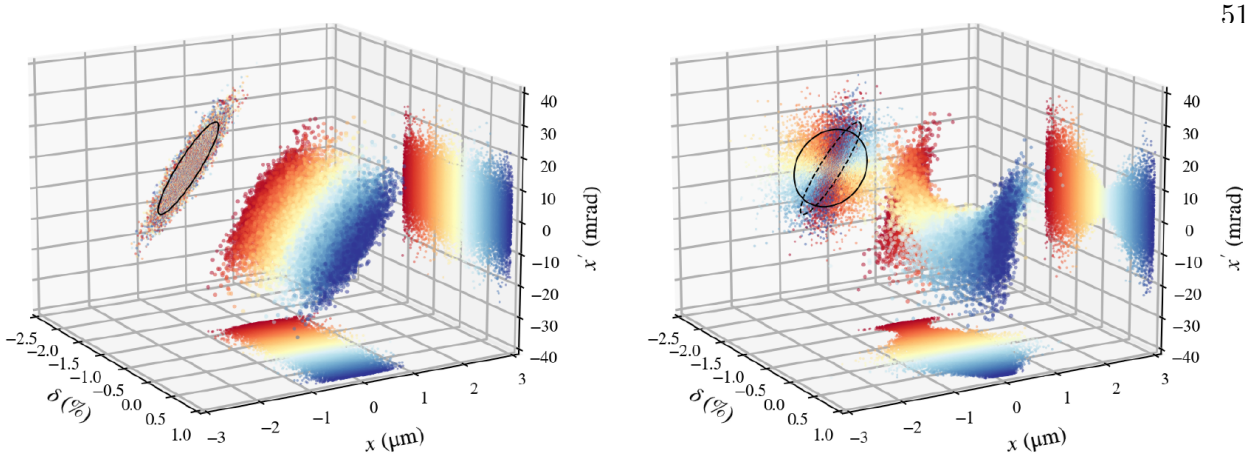

Figure 2.14: Projected emittance of an electron beam before (left) and after (right) propagation through a nonlinear PWFA stage. All energy slices begin with the same CS parameters, but each energy slice rotates at a different frequency while under the influence of the plasma wakefield's transverse focusing force. After the beam has propagated in the plasma long enough, the energy slices have spread out in their rotation and the emittance of the full electron beam has grown larger. Reproduced from Ref. [33]

where β_m is the "matching" beta required for emittance preservation. If the beam initially has a matched betafunction in a uniform plasma density, then it will remain matched and its emittance will not grow.

The trouble with the solution of Eqn. 2.145 is that this leads to extremely small spot size requirements for densities relevant to nonlinear PWFA. Furthermore, this matched beta is only valid for a uniform density plasma and does not account for the beam's evolution in density ramps that exist between the main PWFA stage and the vacuum beamline before and after. As such, Eqn. 2.145 is only the matching condition for the start of the bulk PWFA where the density is constant, and the matching condition for the full plasma including ramps is much more difficult to solve for. The matching conditions for these realistic plasma profiles can be solved for numerically (Ref. [57]) or analytically (Ref. [58, 33]), and are of vital importance to designing a working PWFA.

2.4 Numerical Simulations

2.4.1 PIC Simulations

Analytic efforts can only get us so far, so simulations are required to advance further in modeling the complex processes of electron-beam/plasma interactions. The full electromagnetic interaction can be simulated numerically using a particle in cell (PIC) algorithm. Here particles are modelled using "macroparticles," where a single macroparticle can represent many real particles that have similar position in 6D phase space. A macroparticle can also be weighted to reflect different nominal densities between different macroparticles. The phase space of every macroparticle is used to calculate the electromagnetic fields using Maxwell's equations, and these fields are used to push the macroparticles in phase space. This is repeated as a continuous loop with the general procedure:

(1) Calculate charge and current densities from macroparticle distribution onto a grid.

(2) Calculate electromagnetic fields using Maxwell's equation.

(3) Find the fields at each macroparticle position

(4) Update each macroparticle's 6D phase space using Lorentz force law.

In simulating PWFA, the simulation grid exists in a co-moving window traveling along with the electron beams. Instead of requiring the full plasma length for each simulation, we only need to simulate a region that contains the electron beams and the plasma wake. Still, properly modeling PWFA in 3D requires the use of supercomputers and even then it can be a lengthy process. For example, a nonlinear plasma wake at $n_0 = 3.5 \times 10^{16}$ cm^{-3} has a blowout size on the order of the plasma wavelength $\lambda_p \approx 175$ μm. The total length should be twice this distance in all dimensions. The smallest length scale is set by the size of the electron beam, and to properly resolve a $\sigma = 4$ μm beam we would want a grid size of 0.8 μm. The time step is given by the ratio of the grid size to the speed of light through the Courant-Friedrichs-Lewy (CFL) condition. For the PIC simulation

software we use, VSim [59], if there are 8 macroparticles per grid cell for both the plasma and the beam then it takes about 1 second per step using 1600 cores on NERSC's Cori supercomputer. This results in a time requirement of about 3.5 hours per cm of plasma simulated. Short plasma lenses on the scale of a millimeter or less are not terribly time consuming on an individual basis, but detailed parameter scans remain burdensome.

2.4.2 Particle Tracking

While PIC simulations can be great at being able to model everything that is going on, the price for this information is the time it takes to simulate it. On the other hand, if we already know the wakefields either analytically from Sec. 2.3 or numerically using a PIC simulation, then it can be much more effective to quickly propagate a beam through these fields. Particle tracking codes do just this, solving the evolution of a particle beam's 6D phase space using the equations of motion from Sec. 2.1. Particle trackers also use macroparticles in the same way that PIC codes do; a single macroparticle can effectively represent many real particles, reducing the computational cost of simulating the beam as a whole.

In this dissertation, we often employ the use of a particle tracking algorithm[13] to quickly model the evolution of an electron beam in drift space or in a known external electromagnetic field. To model nonlinear plasma wakes in particle tracking algorithms, we assume that the transverse focusing force is given by the linear function of Eqn. 2.143 and this only varies due to the change in the plasma density $n_0(z)$ that the electron beam is passing through. Effectively, we are treating the plasma wake in 1D and ignoring any finer details. Plasma acceleration is also included here, although as a much simpler version of Eqn. 2.139 which assumes a constant E_z across the beam which depends on the plasma density $n_0(z)$. If the nominal density for PWFA acceleration is $n_0(0)$, then E_z can be lower or even negative if $n_0(z) < n_0(0)$ to represent the blowout changing size and the witness beam being in a relatively different phase of the wake. Beam loading is not included in this particle tracking code, for that one really needs to run a full 3D PIC simulation to simulate

[13] The code we use is called WARGSim, and was programmed in Python by our research group.

the electron sheath's behavior.

We can also import macroparticles from PIC simulations into our particle tracking code. This is useful when one needs to perform a PIC simulation to resolve the dynamics of a short length of beam-plasma interaction, but would then want to propagate the beam afterwards through meters of drift space and conventional magnetic optics. This technique is used in Chapters 3 and 4 to study an electron beam's focus after a PIC simulation of a plasma lens.

2.4.3 Fluid Simulations

Lastly, in considering the density profile of the plasma for PWFA and plasma lenses we need to know what the density profile of the pre-ionized gas is. In the case we are using a gas delivery system, such as a gas jet, it is helpful to numerically simulate the fluid dynamics of a gas passing through the gas jet nozzle. For our plasma lens studies, we use the open-source, computational fluid dynamics software, OpenFOAM [60], to simulate a compressible gas propagating through nozzles and into a vacuum chamber. This allows us to have a 3D density profile of the neutral gas and simulate how the nonlinear plasma wake regime behaves when the plasma density is not uniform.

OpenFOAM itself is a collection of object-oriented classes. This allows for the simulation of fluids in different regimes under the same computational framework. For the regime of a high-pressure nozzle expelling gas into a vacuum, we use the "sonicFoam" solver which is transient[14] , compressible, and allows for the inclusion of turbulence. For our simulations we included the RAS turbulence model, "kOmegaSST," although turbulence was not too much of an issue for a single gas jet outflow into vacuum.

Fluid simulations involve solving the Navier-Stokes differential equations that govern the motion and distribution of fluids. This also includes boundary conditions to specify features in the simulation domain. This includes the walls of the gas jet nozzle, which are modeled to have a "zeroGradient" in pressure but transmissive to temperature changes. We also set the inlet of the

[14] A transient solver iterates the fluid density across each timestep to find how it changes over time. A steady-state solver would just calculate the final, steady-state of the system.

nozzle to have a high, constant pressure to simulate a reservoir, and the low-density region beyond the exit of the nozzle is bounded by boundaries which allow fluid to leave the simulation domain. Due to the extreme density gradient close to the nozzle inlet in the first timesteps, we simulate the early evolution with a much smaller timestep to resolve this gradient.

The time scale of gas motion in a gas jet outflow ($\sim \mu$s) is much longer than the time scale of the plasma ionization and electron beam interaction processes. So, when considering the gas density for these further studies, we assume that the neutral gas is stationary with a density given by a steady state approached by the OpenFOAM algorithm.

Chapter 3

Beam Dynamics in a Passive Plasma Lens

As with any focusing optic, the two components to consider when calculating a focal length are the focusing strength and the thickness of the optic. The focusing strength of a plasma lens arises from the strength of the electromagnetic field within the plasma wake produced by a driver. In Chapter 2 we introduced how a nonlinear blowout wake is generated and derived the electromagnetic fields within the wake. There is a wealth of research investigating the fields within plasma wakes for the purposes of plasma wakefield acceleration, and the passive plasma lens is, in simplified terms, a short plasma wakefield accelerator. Due to the short longitudinal length scales of passive plasma lenses, the energy gain/loss is negligible and we only care about the transverse forces in these wakes. And, as with plasma acceleration, the form of the plasma focusing force depends on if the wake is in the linear or nonlinear regime.

First, we briefly look at the overdense regime and catalog its drawbacks. Then we move on to the nonlinear blowout regime and examine the transverse wakefields of this regime in greater detail. Next, we investigate some of the constraints for this regime and sources of focusing aberrations. We examine some hypothetical use cases for these passive plasma lenses for matching beams into plasma wakefield accelerators for emittance preservation and for strongly focusing beams to emit hard synchrotron radiation. Lastly, we use numerical PIC simulations to verify analytic theory and explore more intricate variations of the passive plasma lens.

3.1 Plasma Lenses in the Linear Overdense Regime

Assuming an electron beam driver in a plasma, the condition for the linear, overdense regime is that the plasma number density is larger than the peak drive beam density [61]. As we saw in Sec. 2.3.1, the linear plasma wake develops in two separate ways depending on the relative length of the drive beam.

If the drive beam is very short, then the plasma density perturbation is a sinusoidal function (Eqn. 2.92 from Sec. 2.3.1) and the resulting wakefields are given by Eqns. 2.99 and 2.100. We would then consider placing a second witness beam at a location in this linear wake where the transverse electric field provides a focusing force. However, as we discussed in Sec. 2.3.2, this focusing force is not a linear function in r and so it will introduce aberrations on the beam.

The other way to design an overdense plasma lens is to use a single, long drive beam that is of relatively low current. In this regime, the electron beam acts as a perturbation on the plasma and the plasma electrons move in a linear response to cancel out the repulsive space-charge field of the electron beam (Fig. 2.6(a)). If the electron beam is much longer than the plasma wavelength and the beam size is much smaller than the plasma wavelength, or $\sigma_z \gg c/\omega_p \gg \sigma_r$, the space-charge field of the beam is perfectly canceled out and the attractive self-generated azimuthal magnetic field $B_\theta = 2\pi n_b e r$, provides the focusing force.[1] Assuming the electron beam has a uniform density ${n_b}^2$, the radial Lorentz force is

$$F_{r,lin}/r \approx 2\pi n_b e^2 \approx 3 \times 10^{-9} n_b[\mathrm{cm}^{-3}]\ \mathrm{G/cm}. \tag{3.1}$$

Both of these schemes in the overdense regime are characterized by a focusing force dependence on the electron beam's density profile, $n_b \equiv \rho_\parallel(\xi)\rho_\perp(r)$, and it is only a linear focusing force if the electron beam's density is very carefully designed. For realistic electron beam density profiles, the focusing force will exhibit aberrations that are detrimental to the quality of the electron beam focus (Fig. 3.1).

[1] See Sec. 2.1.6 for a derivation of the self-fields of a charged particle beam.
[2] A bold assumption.

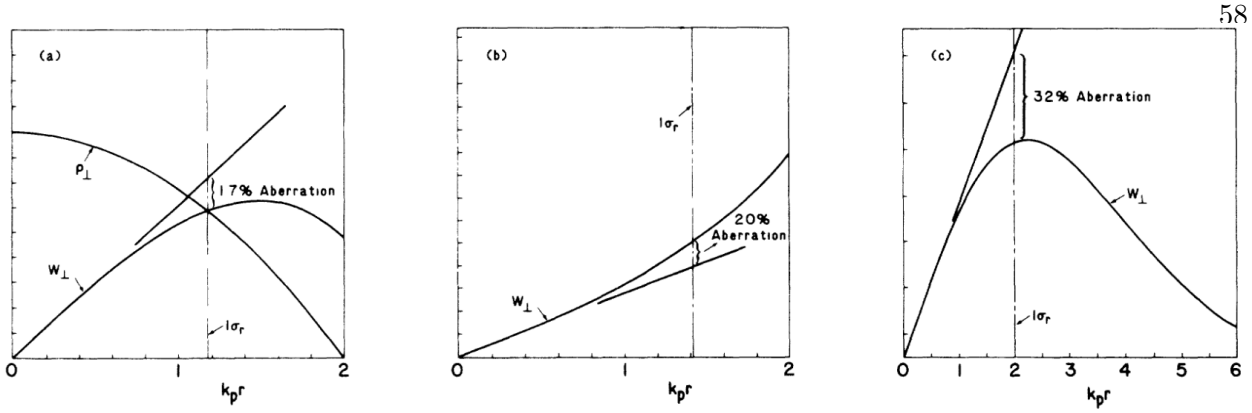

Figure 3.1: Various examples of the focusing force vs r for different electron beam density profiles for length scale $a = 0.2c/\omega_p$: (a) uniform $\rho_\perp = \rho_0$ for $r \leq a$, (b) parabolic $\rho_\perp = \rho_0(1 - r^2/a^2)$, and (c) Gaussian $\rho_\perp = \rho_0 \exp(-r^2/2a^2)$. Reproduced from Ref. [61]

Despite these drawbacks, there has been a wealth of research in overdense plasma lenses [20, 21, 62]. A strong benefit of the overdense regime is its ability to focus positron beams with the same plasma response as for electrons, but reversed [22]. The requirement for the beam density to be smaller than that of the plasma allows for this regime to be reached more readily than the underdense regime. However, this requirement also means that the next-generation beams that are of a much higher current will be too dense for the overdense regime. For these high-current beams, we will need to operate in the nonlinear blowout regime anyways. But as we will see, this regime is more suitable for focusing without aberrations from the focusing force.

3.2 Plasma Lenses in the Nonlinear Blowout Regime

As previously discussed in Sec. 2.3.3, the requirement for the nonlinear, blowout regime is a peak drive beam density that is approximately larger than the plasma number density [46]. Here, the large fields from the electron beam completely expel the plasma electrons and leave behind a "bubble" of stationary ions. The size and shape of this bubble are largely dependent on the parameters of the drive beam, and the length scale of this blowout bubble is on the order of the plasma wavelength. Surrounding this bubble is an electron sheath with a high density of plasma electrons which were blown out by the drive beam, and for a uniform initial plasma density this

sheath current will be axisymmetric. The focusing force on an electron, then, is from the stationary ions within the bubble.

If we again look at the transverse focusing force within the nonlinear blowout wake from Eqn. 2.143, we can start with the following:

$$F(r) = 2\pi n_p e^2 r, \tag{3.2}$$

which is written in cgs units here.[3] For a uniform plasma density n_p, Eqn. 3.2 is linear and axisymmetric. In order for the plasma lens to be considered "thin", the lens must be significantly shorter than one betatron period along the axial dimension, and the plasma density must be rapidly truncated on its upstream and downstream ends. The plasma is then treated as an axisymmetric focusing lens following the standard formalism, where the transverse electron dynamics in the plasma lens can be represented by Hill's Equation from Sec. 2.1.2:

$$x'' + K_x(z)x = 0. \tag{3.3}$$

Here, x represents the transverse displacement from the center of the ion cavity, z is the beam propagation axis, and K_x is the focusing strength of the plasma given by $K_x(z) = 2\pi r_e \frac{n_p(z)}{\gamma_L}$, where constant $r_e \equiv e^2/(m_e c^2)$ is the classical electron radius in cgs units. The focusing is purely linear in x due to the spatial uniformity of the ion density in the blowout cavity. For a uniform plasma density the same equation would hold for transverse axis y.

Plasma lenses can be modelled using transfer matrix formalism from Sec. 2.1. We can start with the typical 2×2 transfer matrix for an individual electron's transverse position and momentum coordinates, x and x',

$$\begin{pmatrix} x \\ x' \end{pmatrix} = M \begin{pmatrix} x_0 \\ x'_0 \end{pmatrix} \tag{3.4}$$

[3] Cgs units are a bit more popular in plasma physics while SI is typically used in accelerator physics. The "smoking gun" that an electromagnetic equation is in cgs is the inclusion of π instead of ϵ_0. A quick substitution to get from SI to cgs is to let $\epsilon_0 \to 1/(4\pi)$.

and examine the transfer matrix for a focusing optic from Eqn. 2.14:

$$M_{thick} = \begin{pmatrix} \cos(\sqrt{K}L) & \frac{1}{\sqrt{K}}\sin(\sqrt{K}L) \\ -\sqrt{K}\sin(\sqrt{K}L) & \cos(\sqrt{K}L) \end{pmatrix} \tag{3.5}$$

with L as the thickness of the lens. As we saw previously in Sec. 2.1.2, Eqn. 3.5 is general to any focusing optic with a linear focusing force, as is the case for the focusing axis of a quadrupole magnet or a thick plasma lens. As we want to consider a plasma lens in the thin lens regime, we can take the limit where $\sqrt{K}L \ll 1$ and this matrix reduces to

$$M_{thin} = \begin{pmatrix} 1 & 0 \\ -KL & 1 \end{pmatrix}. \tag{3.6}$$

Equation 3.6 can be compared to the typical transfer matrix for thin lenses in optics, where the focal length is defined as the negative inverse of the M_{21} element. The focal length of the thin plasma lens is given by

$$f \equiv \frac{1}{KL} = \frac{1}{2\pi r_e}\frac{\gamma_L}{n_p L} \tag{3.7}$$

as previously noted in Refs. [61, 25], where it is assumed that the betatron phase advance in the lens is small: $\Delta\psi \equiv \sqrt{K}L \ll 1$. For a plasma lens with a larger density and/or thickness, the focal length is shorter as the lens focuses an electron beam more strongly. Additionally, a more energetic electron beam with higher longitudinal momentum will contribute to a larger focal length as that electron beam will be more difficult to focus.

For a non-uniform longitudinal plasma density profile, the effective focal length can be calculated by following the standard compound lens formalism:

$$M_{thin,full} = \prod_{i=1}^{k} M_{thin,i} \quad \rightarrow \quad M_{21} = \sum_i M_{21,i}. \tag{3.8}$$

By making the thicknesses of each plasma slice L_i infinitesimally small, the focal length of the full, thin plasma lens instead is calculated as an integral over the density profile:

$$f = \frac{1}{2\pi r_e}\frac{\gamma_L}{\int n_p(z)dz}. \tag{3.9}$$

This is demonstrated in Fig. 3.2, where a plasma of a more realistic density profile has the same effective focal length as a step-function density profile of equivalent density integral. Therefore, so long as the full plasma lens is considered thin we need not include density variations in modelling beam dynamics in the plasma lens. This is not true for thick plasma lenses, however, as the electron beam's spot size significantly evolves within the plasma lens and the focusing dynamics change.

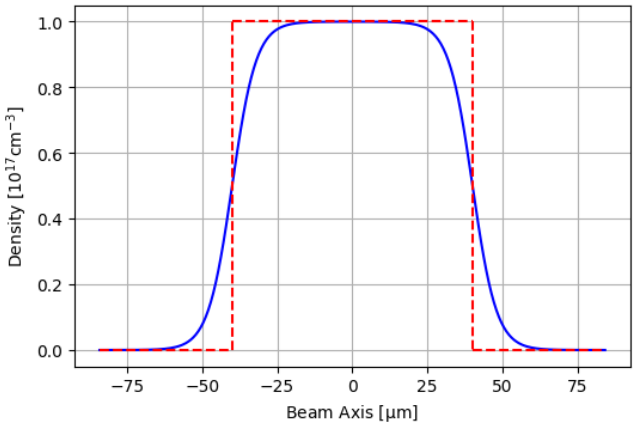

Figure 3.2: Longitudinal plasma density profiles of two thin plasma lenses with identical focal lengths. The integrated plasma density of the curve in solid blue gives a focal length equivalent to the 80 μm thick step-function, and so both plasma lenses can be modelled identically in thin lens formalism.

Here we can consider how the focal length compares between passive thin plasma lenses, conventional quadrupole magnets, and PMQs. We assume an equivalent phase advance through each, $\Delta\psi = \sqrt{K}L = 0.1$, so that all three optics are in the thin regime. The smallest focal length $f = (\sqrt{K}\Delta\psi)^{-1}$ is then dependent on maximizing the focusing strength K. Table 3.1 summarizes the differences in focusing strength K, lens length L, and focal length f for a 10 GeV electron beam focused by a conventional quadrupole magnet with field gradient $G = 1$ T/m, a PMQ with $G = 500$ T/m, and a passive thin plasma lens with density $n_p = 10^{17}$ cm^{-3} corresponding to an equivalent field gradient of 295 kT/m. It can be seen that the laser-ionized, beam-driven, passive thin plasma lens can focus a relativistic electron beam with orders of magnitude greater strength than either of the magnetic focusing devices.

Rather than calculating the trajectories of individual electrons using the transfer matrices

Table 3.1: Comparison between focusing elements for a 10 GeV electron beam with equivalent betafunction phase advance $\Delta\psi = \sqrt{K}L = 0.1$. Focusing strength K comes from magnetic field gradient G for magnetic optics, and from plasma density n_p for the underdense plasma lens.

Focusing Element	K [m^{-2}]	L [mm]	f [cm]
Conventional Quadrupole: $G = 1$ T/m	0.3	180	1000
Permanent Magnetic Quadrupole: $G = 500$ T/m	150	8.2	81
Underdense Plasma Lens: $n_p = 10^{17}$ cm^{-3}	88400	0.34	3.3

of Eqns. 3.5 or 3.6, it is useful to instead use the 3×3 transfer matrix for the propagation of the Courant-Snyder parameters: α, β, and γ from Sec. 2.1.3. For a relativistic electron beam this is:

$$\begin{pmatrix} \beta(z) \\ \alpha(z) \\ \gamma(z) \end{pmatrix} = \begin{pmatrix} M_{11}^2 & -2M_{11}M_{12} & M_{12}^2 \\ -M_{11}M_{21} & M_{11}M_{22} + M_{12}M_{21} & -M_{12}M_{22} \\ M_{21}^2 & -2M_{21}M_{22} & M_{22}^2 \end{pmatrix} \begin{pmatrix} \beta_0 \\ \alpha_0 \\ \gamma_0 \end{pmatrix}. \tag{3.10}$$

Here β_0, α_0, and γ_0 are the Courant-Snyder parameters at the entrance to the plasma lens and M is the 2×2 transfer matrix for a plasma lens (Eqn. 3.6 for thin, Eqn. 3.5 for thick) and a drift space of distance z:

$$M = \begin{pmatrix} 1 & z \\ 0 & 1 \end{pmatrix} M_{lens} \tag{3.11}$$

The beta function at the waist of the electron beam after the thin plasma lens, β_L^*, can be calculated using Eqn. 3.10 and written in terms of the plasma lens parameters K and L in addition to the initial Courant-Snyder parameters:

$$\beta_L^* = \frac{1}{K^2L^2\beta_0 + 2KL\alpha_0 + \gamma_0}, \tag{3.12}$$

and the waist location after the lens is:

$$z_w = \frac{KL\beta_0 + \alpha_0 - L\gamma_0}{K^2L^2\beta_0 + 2KL\alpha_0 + \gamma_0}, \tag{3.13}$$

These equations represent the thin lens approximation of the thick lens expressions from transfer matrix formalism. To account for thick plasma lenses, it is also possible to derive these quantities by using Eqn. 3.5 in the definition of M in Eqn. 3.11. For thick plasma lenses, the waist size and

location are given as

$$\widetilde{\beta}^*_{L,\text{thick}} = \frac{1}{\widetilde{\beta}_0 \sin^2 \widetilde{L} + \widetilde{\gamma}_0 \cos^2 \widetilde{L} + \alpha_0 \sin(2\widetilde{L})}, \tag{3.14}$$

and

$$\widetilde{z}_{w,\text{thick}} = \frac{(\widetilde{\beta}_0 - \widetilde{\gamma}_0) \sin \widetilde{L} \cos \widetilde{L} + \alpha_0 \cos(2\widetilde{L})}{\widetilde{\beta}_0 \sin^2 \widetilde{L} + \widetilde{\gamma}_0 \cos^2 \widetilde{L} + \alpha_0 \sin(2\widetilde{L})}, \tag{3.15}$$

where the tildes indicate dimensionless parameters normalized by a factor of $\sqrt{K}$ or $1/\sqrt{K}$ as appropriate. Length scales are normalized such that $\widetilde{\beta}^*_0 = \beta^*_0\sqrt{K}$, $\widetilde{L} = L\sqrt{K}$, and $\widetilde{z}_{w,\text{thick}} = z_{w,\text{thick}}\sqrt{K}$. Inverse-length parameter γ_0 is normalized as $\widetilde{\gamma}_0 = \gamma_0/\sqrt{K}$. The respective equations for thin lenses are found by setting $\widetilde{L} \ll 1$.

The task of the plasma lens is to reduce the beam's final beta function with respect to the natural vacuum waist size that would be achieved in the absence of the plasma lens, β^*_v. It is useful to quantify the beta function magnification factor $\mathcal{M}$ in terms of β^*_v and the incoming beta function at the start of the plasma lens, β_0:

$$\mathcal{M} = \frac{\beta^*_L}{\beta^*_v} = \frac{1}{K^2L^2\beta_0\beta^*_v \pm 2KL\sqrt{\beta_0\beta^*_v}\left(1 - \frac{\beta^*_v}{\beta_0}\right)^{1/2} + 1}, \tag{3.16}$$

where the sign in the denominator is the sign of α_0. Noting that $\beta_0 > \beta^*_v$, and assuming that β^*_v is fixed, it can be seen that maximum demagnification occurs when β_0 is large, which is true for any thin, linear focusing element.

3.2.1 Constraints

There are a few constraints on the plasma lens with respect to the blowout regime, which we characterized in Sec. 2.3.3 as requiring both a drive and a witness bunch with the latter located within the blowout wake. First, the plasma must be sufficiently underdense relative to the drive beam to allow the production of a fully nonlinear blowout wake. Second, the period of the wake must be long enough to contain the witness bunch in the first blowout cavity. In practice, the second requirement will restrict the upper limit of plasma density to the nominal density of the target PWFA plasma stage. Lower densities will also function properly since the witness beam will

be well within the first blowout cavity. Third, the blowout wake must be wide enough to contain the beam. The third constraint requires $k_p\sigma_r < 1$, where σ_r is the rms transverse spot size of the beam [63]. This can be expressed in terms of the dimensionless parameter χ:

$$\chi \equiv K\beta_0^2 < \frac{\beta_0}{2\varepsilon_n}, \tag{3.17}$$

where ε_n is the normalized transverse emittance of the beam, here considered to be symmetric in x and y. If $\beta_0 = \beta_v^*$, the magnification then becomes

$$\mathcal{M} = \frac{1}{(\Delta\psi)^2\chi + 1}. \tag{3.18}$$

It can be seen from Eqn. 3.18 that the demagnification of the thin plasma lens is ultimately limited by the ratio β_0/ε_n through χ.

3.3 Chromatic Emittance Growth

Chromatic aberrations will appear for any beam with finite energy spread, which presents the final limitation to the demagnification of the beam spot size. If we revisit Hill's Equation from Eqn. 3.3 where the solution of a particle's motion due to uniform focusing in $K(z)$ is a simple harmonic oscillator, including energy spread within an electron beam means that electrons with different energies will have different oscillation frequencies. In Sec. 2.3.5 we briefly introduced the concept of chromatic emittance growth for a long PWFA stage, where a beam propagating in a plasma will gradually have its emittance increased to a saturation value. For a thin plasma lens, energy spread within a beam will contribute to a range of focal lengths when considering Eqn. 3.7. This effect is demonstrated in Fig. 3.3.

The transverse spot size of an electron beam also depends on the normalized emittance through the definition of the Courant-Snyder parameter β:

$$\sigma_{x,y} = \sqrt{\frac{(\beta_{x,y})(\epsilon_{N;x,y})}{\gamma_L}} \tag{3.19}$$

For many applications of an electron beam focusing optic, the goal is to either have a beam with the smallest size possible or a beam with a precise spot size. In both cases, being able to predict

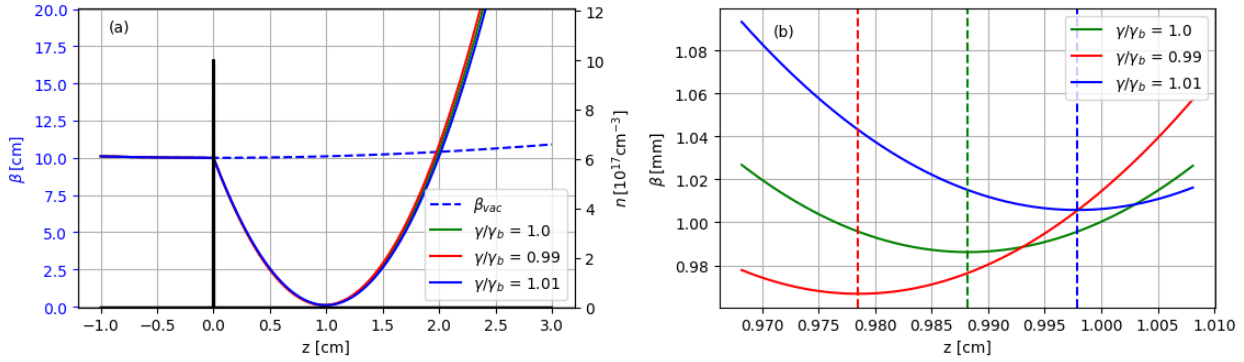

Figure 3.3: Demonstration of chromatic spreading of focal lengths in a plasma lens. Subplot (a) shows the density profile of a plasma lens in solid black with $L = 110.5$ μm, $n_p = 10^{18}$ cm^{-3}, and $f = 1$ cm. The propagation of an electron beam's betafunction before and after the lens is plotted in solid colors, and the unperturbed vacuum propagation is given in dashed blue. The betafunction for an energy slice at the centroid energy, 1% below the centroid energy, and 1% above the centroid is plotted in solid green, red, and blue, respectively. Subplot (b) zooms in around the focus, showing slightly different focus positions (dashed) and sizes for the three energy slices.

the spot size due to chromaticity is crucial for a plasma lens. This is done by including the energy distribution of an electron beam into the modeling of its beta function propagation.

3.3.1 Projected Betafunction

To be able to describe the betafunction for all slices of an electron beam across its energy distribution, we introduce the projected Courant-Snyder parameter:

$$\bar{\beta} \equiv \int_{-\infty}^{\infty} \beta(\delta) f(\delta) d\delta \tag{3.20}$$

with $\beta(\delta)$ as the usual betafunction evaluated at different energy slices of δ, and $f(\delta)$ as a normalized distribution function of the electron beam's energy spectrum. The function $f(\delta)$ can be an arbitrary distribution and Eqn. 3.20 can be evaluated numerically, but for the remainder of this chapter we shall consider two specific cases: the Gaussian energy distribution

$$f(\delta) = \frac{1}{\sqrt{2\pi}\sigma_E} e^{\frac{-\delta^2}{2\sigma_E^2}} \tag{3.21}$$

and the flattop energy distribution

$$f(\delta) = \begin{cases} 1/(2\delta_E), & \text{if } -\delta_E < \delta < \delta_E \\ 0, & \text{elsewhere} \end{cases}, \tag{3.22}$$

where σ_E and δ_E are the spreads in energy for the two distributions, respectively.

Without making further approximations, we can only analytically solve Eqn.3.20 using the formalism for thin lenses in the transfer matrices of Eqn. 3.10 and the flattop energy distribution of Eqn. 3.22. The function $\beta(\delta)$ is found by solving for $\beta(z)$ using the transfer matrix of Eqn. 3.10 at the waist position of the centroid energy slice (Eqn. 3.13), while letting $K(\delta) = K_0/(1+\delta)$. For these calculations, we write the initial Courant-Snyder parameters of the electron beam in terms of the initial vacuum betafunction waist β_i^* and the distance d between this initial waist and the entrance of the plasma lens.

$$\begin{bmatrix} \beta_0 \\ \alpha_0 \\ \gamma_0 \end{bmatrix} = \begin{bmatrix} \beta_i^* + d^2/\beta_i^* \\ -d/\beta_i^* \\ 1/\beta_i^* \end{bmatrix}. \tag{3.23}$$

We also normalize all length scales to $\sqrt{K_0}$ such that $\widetilde{\beta}_i^* = \beta_i^*\sqrt{K_0}$, $\widetilde{L} = L\sqrt{K_0}$, and $\tilde{d} = d\sqrt{K_0}$. With this formalism, we find for the projected betafunction

$$\bar{\beta}\sqrt{K} = \frac{1}{2\widetilde{\beta}_i^*\delta_E\left(1 - 2\widetilde{d}\widetilde{L} + (\widetilde{\beta}_i^{*2} + \widetilde{d}^2)\widetilde{L}^2\right)^2}\left(\frac{2\delta_E}{1-\delta_E^2}(A_{(1)} + B_{(1)}) - C_{(1)}\tan^{-1}(\delta_E)\right); \tag{3.24}$$

with terms $A_{(1)}$, $B_{(1)}$, and $C_{(1)}$ defined

$$A_{(1)} \equiv -\widetilde{\beta}_i^{*2}(\delta_E^2 - 1) + 4\widetilde{\beta}_i^{*2}\widetilde{d}(\delta_E^2 - 1)\widetilde{L} - (\widetilde{\beta}_i^{*2} + \widetilde{d}^2)\left(\widetilde{d}^2(\delta_E^2 - 2) + 3\widetilde{\beta}_i^{*2}(\delta_E^2 - 1)\right)\widetilde{L}^2, \tag{3.25}$$

$$B_{(1)} \equiv 2\widetilde{d}(\widetilde{\beta}_i^{*2} + \widetilde{d}^2)^2(\delta_E^2 - 2)\widetilde{L}^3 - (\widetilde{\beta}_i^{*2} + \widetilde{d}^2)^3(\delta_E^2 - 2)\widetilde{L}^4, \tag{3.26}$$

$$C_{(1)} \equiv 4\widetilde{L}(-\widetilde{d} + (\widetilde{\beta}_i^{*2} + \widetilde{d}^2)\widetilde{L})(\widetilde{\beta}_i^{*2} - \widetilde{d}(\widetilde{\beta}_i^{*2} + \widetilde{d}^2)\widetilde{L} + (\widetilde{\beta}_i^{*2} + \widetilde{d}^2)^2\widetilde{L}^2). \tag{3.27}$$

In Fig. 3.4 we plot the projected spot size calculated using Eqn. 3.20 in both the thin and thick plasma lens regimes, and compare to the spot size statistically calculated from a numerically propagated electron beam with $\beta_i^* = 5\text{cm}$, $\epsilon_N = 3\ \mu\text{m} - \text{rad}$, and with a flattop energy distribution

of $\delta_E = 0.05$. The thin and thick regimes agree well for thicknesses $\sqrt{K}L < 0.3$, and the thick regime is valid for larger thicknesses. The spot size calculated numerically from integrating Eqn. 3.20 for a thick plasma lens matches the spot size for the numerically propagated electron beam. It also matches the spot size calculated from the analytic chromatic amplitude for a thick plasma lens, which is derived from assumptions introduced in the following section.

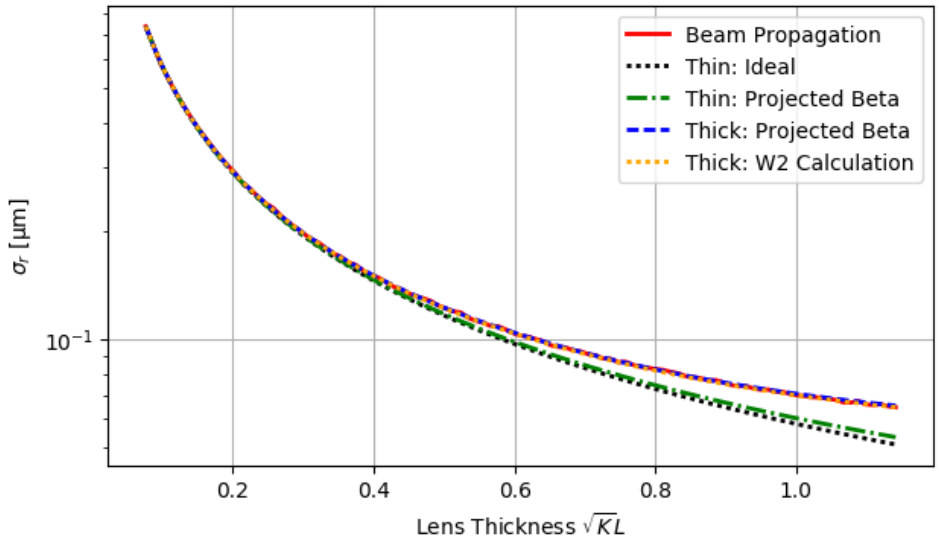

Figure 3.4: Comparison of electron beam spot sizes at a plasma lens focus calculated by different means, varying the plasma lens thickness from the thin to the thick regime. Plasma density $n_p = 10^{18}$ cm^{-3}, initial betafunction $\beta_0 = 5$ cm, initial normalized emittance $\epsilon_N = 3$ μm $-$ rad, and energy spread $\delta_E = 1\%$. Solid red is the true rms spot size from particle tracking simulation. Dotted black is the ideal analytic spot size assuming thin lens and no chromaticity. Dashed green is the analytic spot size including chromaticity from a flattop energy distribution. Dashed blue is the numerically calculated spot size from solving the projected betafunction for a flattop energy distribution and a thick plasma lens. Dotted orange is the analytic solution for a thick plasma lens using the chromatic amplitude and assuming a Gaussian energy distribution. For the case presented here, both the numerically calculated projected betafunction and the chromatic amplitude for a thick lens match the spot size from numerical beam propagation, which is larger than the ideal, thin plasma lens regime.

In the case of a Gaussian energy distribution, approximations can be made which allow for the first-order emittance growth to be calculated. This calculation is performed in Ref. [64], and the general solution is applied to plasma lenses in Sec. 3.3.2.

3.3.2 Chromatic Amplitude

In solving for the projected betafunction for an electron beam with a Gaussian energy distribution, the main approximation used in Ref. [64] is that $\beta(\delta)$ is approximately linear. In the case of a plasma lens this approximation is valid, as shown in Fig. 3.5. This is true for even thick lenses, so long as the energy spread within the beam is not extreme.

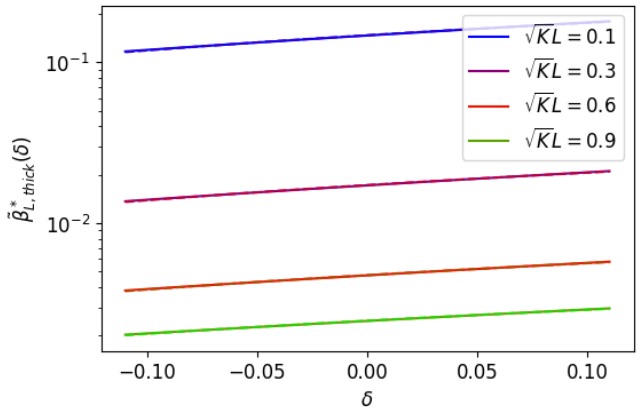

Figure 3.5: Betafunction of an electron beam after a plasma lens for different energy slices with 10% energy variation. Plotted on each curve is a linear fit, and each curve is effectively linear. This leads into the assumption where the chromatic amplitude definition is valid if the betafunction is linear in δ.

With these approximations, a chromatic amplitude is defined in Ref. [64] as:

$$W^2 = \left(\frac{\partial\alpha}{\partial\delta} - \frac{\alpha}{\beta}\frac{\partial\beta}{\partial\delta}\right)^2 + \left(\frac{1}{\beta}\frac{\partial\beta}{\partial\delta}\right)^2 \tag{3.28}$$

where $\partial/\partial\delta$ are derivatives with respect to a relative energy offsets δ centered at $\delta = 0$. This chromatic amplitude is used to calculated the first-order geometric emittance growth. The ratio of the projected geometric emittance at the exit of the lens, ε_f, to the initial geometric emittance, ε_0, is given by

$$\frac{\varepsilon_f}{\varepsilon_0} \approx \sqrt{1 + W^2\sigma_E^2}, \tag{3.29}$$

with a Gaussian rms spread σ_E.

Equation 3.28 is evaluated in the thin lens regime by letting $K(\delta) = K_0/(1+\delta)$ for a thin lens transfer matrix from Eqn. 3.6. The Courant-Snyder parameters are then found in terms of δ

in Eqn. 3.10, and W^2 can then be calculated. The chromatic amplitude reduces to a compact form which depends on the phase advance through the lens and the ratio of the incoming beta function to the initial beam emittance, captured in the parameter χ:

$$W_{thin}^2 = K_0^2 L^2 \beta_0^2 = (\Delta\psi)^2 \chi. \tag{3.30}$$

The expression for geometric emittance growth in a passive thin lens from Eqns. 3.29 and 3.30 is identical to the definition of aberration power from Ref. [65] for a general rms focusing strength variation.

We can also calculate a general expression for the chromatic amplitude using the thick lens transfer matrix of Eqn. 3.5. For this solution, we use the Courant-Snyder parameters defined using β_i^* and d, as in Eqn. 3.23, and normalize length scales to $\sqrt{K_0}$. The calculation process for the chromatic amplitude of a thick plasma lens is identical to deriving Eqn. 3.30, but the solution is much less compact:

$$W_{thick}^2 = \frac{1}{16}\left[\frac{4\left(1-(1+2\tilde{d}\widetilde{L})\cos(2\widetilde{L})+(\tilde{d}+(-1+\widetilde{\beta}_i^{*2}+\tilde{d}^2)\widetilde{L})\sin(2\widetilde{L})\right)^2}{((\widetilde{\beta}_i^{*2}+\tilde{d}^2)\cos^2\widetilde{L}+\sin^2\widetilde{L}+\tilde{d}\sin(2\widetilde{L}))^2} + \frac{1}{\widetilde{\beta}_i^{*2}}\left(A_{(2)}+B_{(2)}\right)^2\right]; \tag{3.31}$$

with

$$A_{(2)} \equiv 2(-1+\widetilde{\beta}_i^{*2}+\tilde{d}^2)\widetilde{L}\cos^2\widetilde{L} - 2(-1+\widetilde{\beta}_i^{*2}+\tilde{d}^2)\widetilde{L}\sin^2\widetilde{L} + (1+\widetilde{\beta}_i^{*2}+\tilde{d}^2+4\tilde{d}\widetilde{L})\sin(2\widetilde{L}) \tag{3.32}$$

and

$$\begin{aligned} B_{(2)} \equiv & \frac{-2\tilde{d}\cos(2\widetilde{L})+(-1+\widetilde{\beta}_i^{*2}+\tilde{d}^2)\sin(2\widetilde{L})}{(\widetilde{\beta}_i^{*2}+\tilde{d}^2)\cos^2\widetilde{L}+\sin^2\widetilde{L}+\tilde{d}\sin(2\widetilde{L})} \\ & \times\left(1-(1+2\tilde{d}\widetilde{L})\cos(2\widetilde{L})+(\tilde{d}+(-1+\widetilde{\beta}_i^{*2}+\tilde{d}^2)\widetilde{L})\sin(2\widetilde{L})\right) \end{aligned} \tag{3.33}$$

A comparison between the emittance growth predicted by Eqns. 3.30 and 3.31 to the emittance growth simulated in a particle tracking code is shown in Fig. 3.6. The thin lens regime agrees well with simulation for when the normalized thickness $\sqrt{K}L$ is less than 0.3, while the general thick lens expression is consistent across the full range.

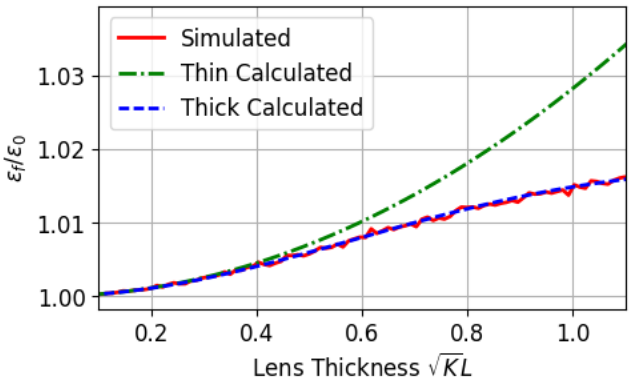

Figure 3.6: Emittance growth (y-axis) versus normalized lens thickness (x-axis) for the analytic expression of Eqn. 3.30 in green, the analytic expression of Eqn. 3.31 in blue, and the emittance growth measured in a numerical particle tracking simulation in red. The jaggedness of the numerical simulation results is due to the limited number of macroparticles in the simulated beam.

3.3.3 Emittance Growth in Drift Space

In addition to the geometric emittance growth within the plasma, the normalized emittance can increase rapidly for a tightly focused beam propagating in vacuum. This normalized emittance growth is quantified in Ref. [56] as

$$\epsilon_n^2 = \langle\gamma_b\rangle^2 \left(s^2\sigma_E^2\sigma_{x'}^4 + \epsilon^2\right) \tag{3.34}$$

where s is the drift distance and $\sigma_{x'} = \sqrt{\gamma\epsilon}$ is the beam's rms size in momentum space. The normalized emittance growth at the focus is written

$$\epsilon_n^2 = \langle\gamma_L\rangle^2 \left((KL\beta_0 + \alpha_0 - L\gamma_0)^2\sigma_E^2 + 1\right)\epsilon_f^2. \tag{3.35}$$

For a passive thin plasma lens of $f = 3.3$ cm focusing a beam of $\beta_v^* = \beta_0 = 5$ cm and $\sigma_E = 0.25\%$, Eqn. 3.29 predicts the geometric emittance to grow by a factor of 7×10^{-6} between the lens and the focus. Equation 3.35 also predicts the normalized emittance to grow further by a factor 7×10^{-6}. Under most realistic conditions, $(KL\beta_0 + \alpha_0 - L\gamma_0)^2$ will be on the order of 1 and σ_E^2 will be on the order of 10^{-4} or smaller, thus the chromatic emittance growth from the passive thin plasma lens will be negligibly small.

3.4 Matching into Plasma Wakefield Accelerators

Here we can examine the effectiveness of a plasma lens for one potential application: that of matching into a plasma wakefield accelerator (PWFA). As introduced in Sec. 2.3.5, to avoid unnecessary emittance growth in a PWFA it is important to focus the incoming electron beam down to a particular transverse beam size that is matched to the focusing strength of the plasma. When a beam is matched, even if the different energy slices in the beam rotate in phase space at different frequencies the transverse size of the full beam's distribution remains constant. Perfect matching can lead to zero chromatic emittance growth under ideal conditions, as discussed in Refs. [57, 58].

Matching can be accomplished using the plasma density transition between vacuum and the PWFA itself [58], although this requires the incoming electron beam to have a particular vacuum betafunction waist determined by the density ramp. This additionally requires control over the density ramp profile, which is not a trivial process. Alternatively, rather than relying on the density ramp of the PWFA we can instead consider a thin plasma lens upstream of the PWFA that produces the necessary focusing to correctly reach the matching condition. This section seeks to look at realistic examples to see if a plasma lens could feasibly be used to match an electron beam into a PWFA without significantly increasing the emittance.

There are a number of ways to approach PWFA matching, but here we can assume a typical case where we would know the plasma density of the PWFA, thus giving the matching betafunction waist and position, as well as the beam parameters of the incoming electron beam. We are then free to choose the position and focal length of a plasma lens that would satisfy matching. Lastly, from constraints on the plasma lens density one can determine what the appropriate plasma lens density and thickness are from the focal length.

To solve for the plasma lens focal length and position, we first take Eqn. 3.12 and Eqn. 3.13. Using Eqn. 3.23 we re-express in terms of the initial vacuum betafunction waist β_i^* and distance

from plasma lens d:

$$\beta_f^* = \frac{\beta_i^*}{(KL\beta_i^*)^2 + (1 - KLd)^2} \tag{3.36}$$

and

$$z_{\text{waist}}^* = \frac{(KL)(\beta_i^{*2} + d^2) - d}{(KL\beta_i^*)^2 + (1 - KLd)^2}. \tag{3.37}$$

The waist betafunction after the plasma lens is set to the matching condition $\beta_f^* = \beta_{\text{match}}^*$. We also define the length z_m as the distance between the vacuum waist positions of the initial and final betafunctions. If we set the coordinates of the z axis such that $z = 0$ corresponds to the entrance of the plasma lens, then $z_m \equiv z_{\text{waist}}^* + d$. A positive d corresponds to the initial waist being upstream of the plasma lens, while z_{waist}^* is always positive.

We then reorganize these equations to solve for the two unknowns of inverse focal length KL and the position of the plasma lens with respect to initial vacuum waist d:

$$KL = \frac{2\beta_i^*\beta_{\text{match}}^* z_m + (\beta_i^* + \beta_{\text{match}}^*)\sqrt{\beta_i^*\beta_{\text{match}}^*\left(\left(\beta_i^* - \beta_{\text{match}}^*\right)^2 + z_m^2\right)}}{\beta_i^*\beta_{\text{match}}^*\left(\left(\beta_i^* + \beta_{\text{match}}^*\right)^2 + z_m^2\right)} \tag{3.38}$$

and

$$d = \frac{\beta_i^* z_m - \sqrt{\beta_i^*\beta_{\text{match}}^*\left(\left(\beta_i^* - \beta_{\text{match}}^*\right)^2 + z_m^2\right)}}{\beta_i^* - \beta_{\text{match}}^*}. \tag{3.39}$$

These equations are generally applicable for finding a plasma lens for PWFA matching, so long as the initial and matching vacuum betafunctions intersect upstream of the matching betafunction's waist. There also exists a disconinuity where the initial focus is equivalent to the matching focus, $\beta_i^* = \beta_{\text{match}}^*$, and Eqn. 3.39 instead is solved as

$$d = z_m/2 \tag{3.40}$$

In this special case, the plasma lens is placed halfway between the initial and matching focus positions. However, this is only valid if the initial waist is upstream of the matching waist, or $z_m < 0$.

3.4.1 Plasma Lens at Initial Betafunction Waist

First, we examine a hypothetical PWFA matching scenario with a plasma lens placed at the incoming beam's vacuum betafunction waist that needs to halve the betafunction to achieve matching [66]. We consider a PWFA with a 10 cm-long uniform density flat-top region of 3×10^{16} cm^{-3} and Gaussian density ramps on either side with a half-width at half-maximum of $\sigma_{hw} = 2.54$ cm. For a 10 GeV electron beam to be matched into this plasma source, it must have a vacuum waist beta function of $\beta^* = 2.5$ cm at a position 4.55 cm before the start of the flattop [57].

The smallest vacuum beta function that can be produced by the final focus quadrupole electromagnets at the upcoming FACET-II facility is $\beta_v^* = 5$ cm [67, 11]. An appropriate thin plasma lens can provide the additional focusing necessary to reach the matching beta function. By solving Eqn. 3.18 with $\mathcal{M} = 1/2$ and $n_p = 3 \times 10^{16}$ cm^{-3}, the required lens thickness is found to be 737 μm. The ideal location of the plasma lens can be inferred from Eqn. 3.13, which gives the distance from the plasma lens to the new waist location.

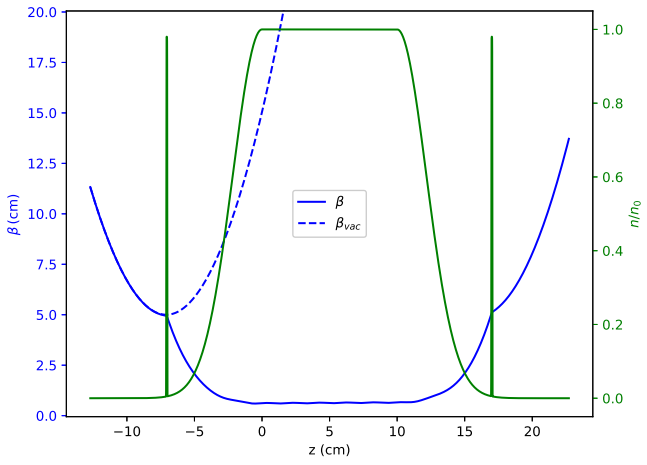

Figure 3.7: Evolution of the beta function of a 10 GeV electron beam through a plasma wakefield accelerator (PWFA) with density ramps of half-width 2.54 cm and peak density $n_0 = 3 \times 10^{16}$cm^{-3}. A thin plasma lens is located at the entrance and exit of the PWFA to match the incoming beam and mitigate its final divergence, respectively. The plasma density profile n_p is given by the solid green line. The electron beam's beta function β is given by the solid blue line. The beta function corresponding to vacuum propagation is given by the dashed blue line.

Figure 3.7 shows the beta function evolution of a 10 GeV electron beam as it propagates

through the thin plasma lens and PWFA described above. The simulation treats the plasma as an axisymmetric, linear focusing element. Fig. 3.7 also includes an identical plasma lens at the PWFA exit that acts to control the high-divergence outgoing beam. After the exit lens, the beam evolves as if propagating from a vacuum waist beta function of $\beta_f^* \simeq 5$ cm. Energy gain in the PWFA is modeled according to Ref. [57], and an increase of 1.72 GeV for the witness beam slightly diminishes the focusing strength of the exit lens with respect to its counterpart at the PWFA entrance. This can be compensated for by simply adjusting the thickness or position of the exit lens to achieve the desired value of β_f^*.

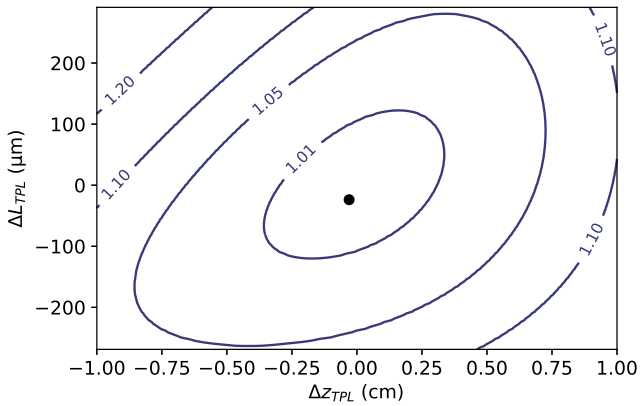

Figure 3.8: Contour plot of the emittance growth $\epsilon_{N,f}/\epsilon_{N,0}$ over a range of errors in the plasma lens thickness ΔL_{TPL} and plasma lens location Δz_{TPL}. A tolerance of 1% saturated emittance growth, corresponding to the $\epsilon_{N,f}/\epsilon_{N,0} = 1.01$ contour, permits an error of up to $\pm 100\,\mu$m and $\pm 0.3\,$cm in the lens thickness and position, respectively. The filled circle represents the parameter set of the thin lens used in Fig. 3.7 with $\epsilon_{N,f}/\epsilon_{N,0} = 1.0004$ and the axis origin represents the perfect matching value.

We also investigate the experimental tolerance on plasma lens thickness and location. Figure 3.8 shows the ratio of the saturated emittance in the PWFA to the initial emittance for various combinations of error in the lens thickness (ΔL_{TPL}) and lens location (Δz_{TPL}). In the FACET-II example considered here, the tolerance in both parameters is large. The lens thickness can vary by $\pm 100\,\mu$m (13.6%) and the lens location can vary by $\pm 0.3\,$cm while still preserving the beam emittance in the PWFA to better than 1%.

The designed lens thickness and position are shown in Fig. 3.8 with the filled circle, and the

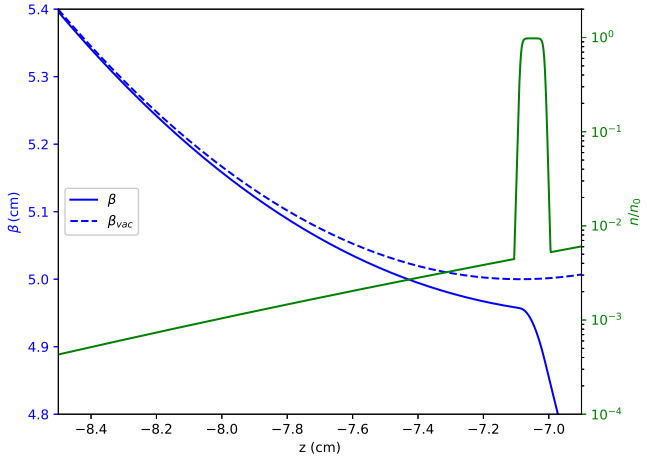

Figure 3.9: Evolution of the electron beam beta function in the region of the plasma lens. Perturbative focusing from the low density plasma ramp shifts the beta function (solid blue) away from its theoretical vacuum trajectory (dashed blue). The plasma density profile n_p is given in logarithmic scale by the solid green line, where $n_0 = 3 \times 10^{16}$ cm^{-3}.

values is slightly different from the perfect matching condition (the axis origin) for two reasons. The first is because we used the thin lens approximation. If instead Eqns. 3.14 and 3.15 were used to design the lens, it would lead to a lens thickness 7.32 μm greater than the design value. The second effect originates from the perturbative focusing of the electron beam by the long tail of the PWFA plasma ramp prior to its arrival at the lens [58]. Fig. 3.9 shows that the result of this effect is to create a beta function at the lens which is slightly smaller than the predicted vacuum value. If necessary, both of the above effects can be analytically treated to generate a more accurate set of lens parameters for perfect beam matching, though as is demonstrated in Fig. 3.8, the simpler thin lens treatment should be more than sufficient for most purposes.

3.4.2 Plasma Lens at Arbitrary Distance from Initial Betafunction Waist

In this second example of PWFA matching with a plasma lens, we look at a more general case in which the plasma lens does not necessarily need to be at the initial vacuum betafunction waist position.

This hypothetical scenario is presented in Fig. 3.10. For the plasma density profile of the

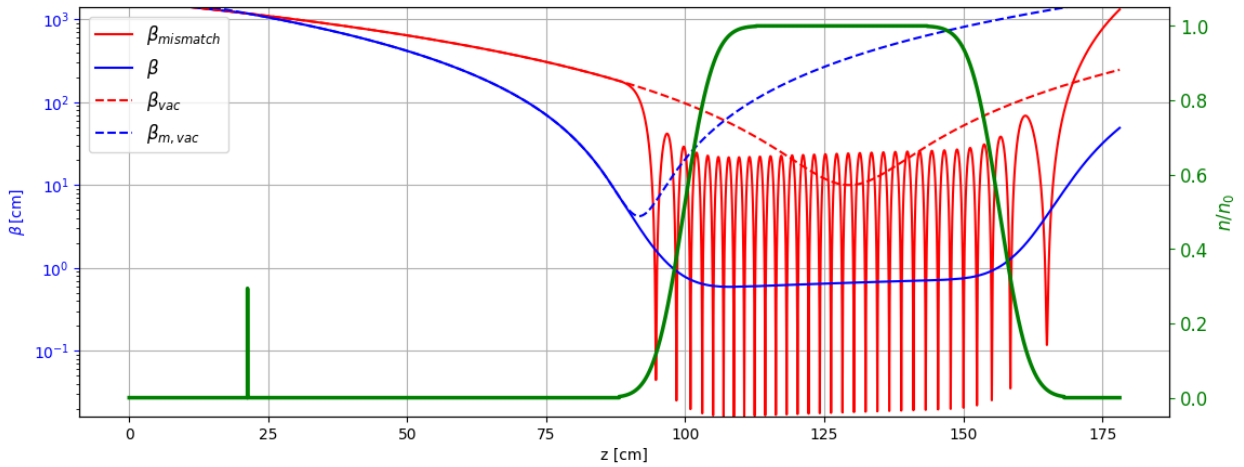

Figure 3.10: Evolution of a beam's betafunction through a longitudinal plasma density profile (green) both using (blue) and ignoring (red) an upstream plasma lens (green, left). The vacuum propagation of the matched (blue) and mismatched (red) beams are shown in dashed lines. Due to the thin plasma lens upstream of the PWFA, the beam receives an appropriate focusing kick to achieve matching.

PWFA, we take a density profile typical of the Lithium oven used at FACET with a peak density of 3.4×10^{16} cm^{-3}. The density ramps are fitted to error functions, which result in a flattop region of 30.8 cm in the middle of the PWFA and a smooth density transition between vacuum and peak density of 24.6 cm. More importantly, the particular density profile given here results in a matching condition for an incoming electron beam to have a vacuum betafunction waist size of $\beta^*_{i,\text{match}} = 4.2$ cm at position $z^*_{i,\text{match}} = 37.8$ cm upstream of the center of the PWFA.

To achieve matching, a thin plasma lens needs to be placed at the intersection of the matched and unmatched vacuum betafunction propagation curves. This is also obtained from Eqn. 3.39, which places a plasma lens 108.2 cm upstream of the initial vacuum waist. Solving Eqn. 3.38 gives the necessary focal length, from which we can choose the plasma lens density and thickness. For a plasma lens far upstream of the waist, we will want a lower plasma density so that the blowout wake is large enough to contain the unfocused electron beam. In Fig. 3.10, we have chosen a plasma density of 1.0×10^{16} cm^{-3}, which leads to a required thickness of 55.1 μm.

This example in particular is quite promising for PWFA matching in experiments at FACET-II. Due to the infrastructure of the facility, it can be quite difficult to generate a plasma lens close to

the plasma source for the PWFA. However, being able to place a plasma lens more than 0.5 m away is much easier to accommodate. Additionally, compared to the example in the previous section, the small thickness of the plasma lens here (55.1 μm compared to 737 μm) can be easier to generate. We will discuss methods and limitations of generating thin plasma lenses in greater detail in the next chapter.

These two examples showcase different approaches to using plasma lenses to achieve matching in a PWFA, although many other configurations are possible. Under ideal operation, the plasma lenses can achieve matching for a wide variety of scenarios by solving Eqns. 3.38 and 3.39, and if the lens thickness is restricted to the thin regime the chromatic emittance growth in Eqn. 3.29 is small as well.

3.5 Hard Synchrotron Radiation and the Oide Limit

In this section we revisit the goal of achieving small spot sizes for high luminosity in colliders. In addition to the chromatic aberrations discussed in Sec. 3.3 for strong plasma lenses, we also need to consider how the emission of synchrotron radiation impacts the electron beam's size at its focus. We briefly introduced the concept of synchrotron radiation in Sec. 2.1.1 to motivate why high-energy electron accelerators are built as linear machines rather than circular. Strong focusing optics can also present a significant curvature to the orbit trajectory of an electron beam, and in return emit hard synchrotron radiation that impacts the phase space distribution of the electron beam.

This process of hard synchrotron radiation was quantified by Ref. [68] in 1988 for beams with a Gaussian distribution. The spot size at the electron beam's post-lens focus is calculated by integrating over the average number of photons emitted by each electron and its corresponding deviation from the ideal trajectory over the entire beam:

$$\sigma_y^{*2} = \beta_y^* \epsilon_y + \frac{110}{3\sqrt{6\pi}} r_e \lambda\!\!\!^{-}_e \gamma^5 F(\sqrt{K}L, \sqrt{K}, l^*) \left(\frac{\epsilon_y}{\beta_y^*}\right)^{5/2} \tag{3.41}$$

where r_e is the classical electron radius, $\lambda\!\!\!^{-}_e$ is the Compton wavelength, and the dimensionless

function F is defined

$$F(\sqrt{K}L, \sqrt{K}l^*) \equiv \int_0^{\sqrt{K}L} \left|\sin\phi + \sqrt{K}l^*\cos\phi\right|^3 \left[\int_0^{\phi} \left(\sin\phi' + \sqrt{K}l^*\cos\phi'\right)^2 d\phi'\right]^2 d\phi \qquad (3.42)$$

with length l^* as the distance between the downstream edge of the focusing optic and the electron beam focus. Equation 3.41 is true for any kind of focusing optic.

An interesting aspect of Eqn. 3.41 is that for any given focusing optic where K and L are fixed, there is a particular β_y^* that minimizes σ_y^*. This minimum spot size is found to be

$$\sigma_{y,min}^* = \left(\frac{7}{5}\right)^{1/2} \left[\frac{275}{3\sqrt{6\pi}} r_e \lambda\!\!\!^{-}_e F(\sqrt{K}L, \sqrt{K}l^*)\right]^{1/7} (\epsilon_{Ny})^{5/7} \qquad (3.43)$$

when the electron beam's waist betafunction is chosen to be

$$\beta_{y,min}^* = \left[\frac{275}{3\sqrt{6\pi}} r_e \lambda\!\!\!^{-}_e F(\sqrt{K}L, \sqrt{K}l^*)\right]^{2/7} \gamma\,(\epsilon_{Ny})^{3/7}\,. \qquad (3.44)$$

Equation 3.43 is commonly referred to as the "Oide Limit", and it is important to consider when designing future linear colliders to determine the maximum possible luminosity given a final focusing scheme. However, there has yet to be experimental verification of the Oide effect due to the high electron beam energies and focusing optic strengths required. Underdense plasma lenses present a unique opportunity here, as the high focusing fields present in the blowout wake can allow for electron beams with currently-available energies to reach the Oide limit with the correct initial parameters. An example of calculations using these equations for the case of strong, underdense plasma lenses is presented in Sec. 3.5.1. While the analytic theory in Ref. [68] shows that the focused electron beam has an increased rms spot size due to synchrotron radiation. It is also the case that the beam's tranverse profile no longer follows a Gaussian distribution. It is shown in Ref. [69] that a more rigorous approach to the beam's distribution reveals a non-Gaussian distribution more favorable for luminosity. The luminosity from this approach is higher than would be expected from just using Eqn. 3.41, but still less than the luminosity from ideal focusing.

A method to avoid significant focusing aberrations due to hard synchrotron radiation is discussed in Ref. [70] whereby a long plasma is used to adiabatically focus the electron beam. The idea of an adiabatic plasma focuser is to have the plasma density gradually increase along the

longitudinal electron beam axis so that the betafunction of the beam decreases gradually towards the interaction point. As a result, electrons that would experience different focusing forces due to chromaticity or synchrotron radiation emission instead are all enclosed within an ever-decreasing beam envelope. This allows for the Oide limit to be surpassed, although there are a few outstanding questions on how to implement such an adiabatic focuser in practice. The most important is that of generating such a long plasma source with a tapered density profile.

3.5.1 Plasma Lens Example of Reaching Oide Limit

In this section we perform a simple calculation of the equations presented in the preceding section for the case of a plasma lens implemented for an electron beam with parameters similar to those expected at FACET-II [11]. The parameters assumed in the following calculations are summarized in Table 3.2.

Table 3.2: Electron beam and plasma lens parameters used in example calculations for hard synchrotron radiation aberrations and reaching the Oide limit with a FACET-II -like beam.

Parameter	Symbol	Value
Electron Beam Energy	E	10 GeV
Lorentz Value	γ	19569.5
Initial Waist Betafunction	β_i^*	500 cm
Normalized Transverse Emittance	ϵ_N	$3\ \mu\text{m} - \text{rad}$
Gaussian Energy Spread	σ_E	0.1%
Electron Beam Charge	Q	1.5 nC
Longitudinal Beam Size	σ_z	$5.2\ \mu$m
Plasma Lens Density	n_0	$10^{18}\ \text{cm}^{-3}$
Plasma Lens Thickness Range	L	$10 - 200\ \mu$m

Using these parameters, we evaluate the equations in Sec. 3.5 and allow for the thickness of the plasma lens to be varied from 10 to 200 μm. Figure 3.11 plots the dimensionless F function of Eqn. 3.42, and the optimal betafunction at the focus (Eqn. 3.44) is plotted in Fig. 3.12. Also plotted in Fig. 3.12 is the betafunction at the focus of an ideal plasma lens (Eqn. 3.14). When the betafunction from the plasma lens matches the optimal betafunction of Eqn. 3.44, then the rms spot size at the focus will equal Eqn. 3.43. Otherwise, the rms spot size is given by Eqn. 3.41.

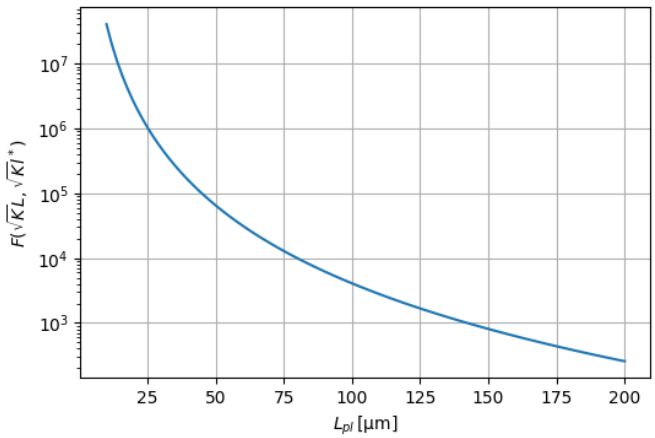

Figure 3.11: Dimensionless F function of Eqn. 3.42 plotted against plasma lens thickness using the electron beam and plasma parameters in Table 3.2.

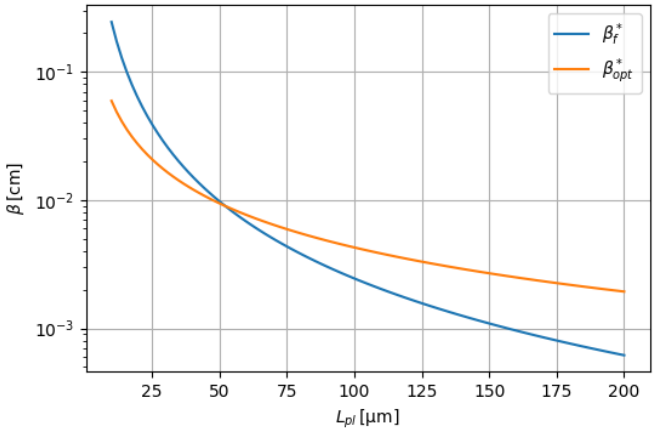

Figure 3.12: Optimal betafunction β^*_{opt} for reaching the Oide limit (solid orange) plotted for a plasma lens with varying thickness. Electron beam and plasma parameters used are from Table 3.2. The ideal betafunction of a thick plasma lens (solid blue) matches the optimal betafunction when the Oide limit is reached for a plasma lens of a set thickness.

The rms spot sizes from Eqn. 3.41 and Eqn. 3.43 are plotted in Fig. 3.13, alongside the spot size from ideal plasma lens focusing and the spot size assuming the only aberration is the chromatic phase spreading described in Sec. 3.3. Assuming that one has the capability to accurately vary the plasma lens thickness, Fig. 3.13 shows a favorable prediction of spot sizes with respect to performing an experiment. Varying thickness allows us to easily reach regimes before and after synchrotron radiation dominates beam aberrations. Keeping the density constant means that the

blowout dynamics will not change as the thickness is varied, so conditions will remain steady across the parameter scan. It is crucial, however, that the initial energy spread of the electron beam is low so that the aberrations from chromaticity do not compete with synchrotron radiation effects.

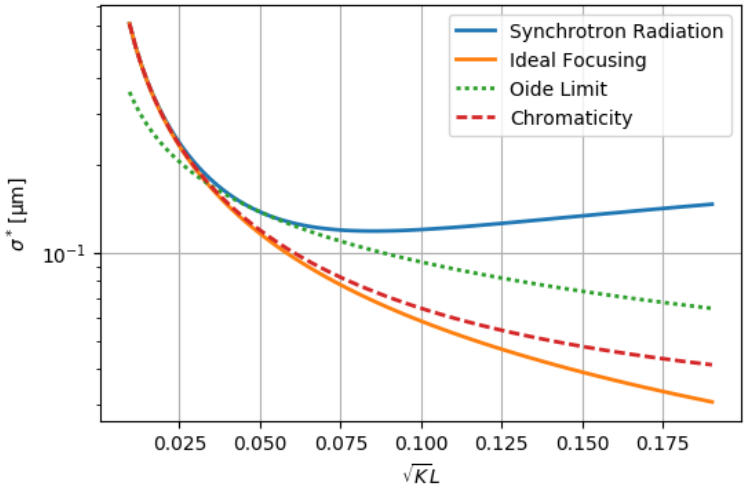

Figure 3.13: Spot size calculations for a varying plasma lens thickness using parameters from Table 3.2. The analytic spot size from synchrotron radiation effects is shown in solid blue, while the Oide limit is plotted in dashed green. As the thickness increases, the Oide limit also decreases. However, from Fig. 3.12 there is only one thickness where the focused betafunction is equal to the optimal betafunction, and in this figure that corresponds to when the dashed green is tangent to the solid blue. Ideal spot size at the focus is plotted in solid orange, and the spot size due to only chromatic effects is plotted is dashed red.

The most challenging parameters presented in Table 3.2 is that of a 500 cm betafunction electron beam fitting within the small blowout wake in a plasma lens of density 10^{18} cm^{-3}. However, in principle there are ways to go about making this experimentally feasible. One could use a very strong driver to blow out an especially wide wake, and one could additionally consider placing the witness bunch within the second wake period (Sec. 3.6.2. In addition, it is not trivial to measure electron beam spot sizes on the order of 100 nm. Current state-of-the-art techniques using laser wirescanners[4] and optical transition radiation (OTR)[5] monitors can reach sub-micron precision but are challenging to implement [72, 71]. Additionally, it may be possible to indirectly measure

[4] A process using Compton scattering of the electron beam off of a laser focus.

[5] From Ref. [71], OTR light appears "when a charged particle crosses a boundary between two media with different dielectric properties," such as placing a metal foil in the path of an electron beamline.

the beam focus size by instead measuring the beam size downstream of the focus.

3.6 Numerical PIC Studies

While analytic modeling can take us quite far in predicting the evolution of an electron beam during and after a thin plasma lens, the full electron beam-plasma interaction dynamics are a complex process that need a 3D PIC simulation to fully model. As described in Sec. 2.4.1, the PIC simulation software we use is VSim. With VSim, we can propagate an electron beam with either one or multiple bunches through a plasma lens. Afterwards, we can export the 6D phase space of the focused electron bunch and perform a simple macro-particle propagation through free space to the focus. This second simulation is considered a particle tracking simulation (Sec. 2.4.2) and simply evolves the phase space of an electron beam using transfer matrix formalism (Sec. 2.1.2).

We also take a look at two theoretical implementations of plasma lenses that are difficult to model analytically. While topics use idealized plasma sources for the plasma lenses, the beam-plasma wake dynamics are quite messy and depend heavily on the particular parameters used. As such, we can only study these qualitatively. However, they can be useful regimes for early plasma lens experiments with limited accessible parameters.

3.6.1 Ideal Two-Bunch Underdense Plasma Lens Focusing

First we look at the propagation of a two bunch electron beam with FACET-II parameters (Table 3.3) through the plasma lens discussed in Sec. 3.4.1. Without considering the PWFA itself, the plasma lens is designed to reduce a $\beta_i^* = 5$ cm electron beam down to $\beta_f^* = 2.5$ cm without introducing any observable focusing aberrations into the beam.

The 3D PIC simulation in VSim was performed using a moving window that is 300 μm wide in the transverse dimension and 202 μm long in the longitudinal dimension. The number of cells are 518 and 360 for the transverse and longitudinal dimensions, respectively. Particles are loaded at the front of the moving window, before the drive beam. The transverse boundaries of the window have a matched absorbing layer which is 8 cells deep. The simulation starts 700 μm before the center

Table 3.3: Electron beam parameters used in PIC simulation based on projected parameters for FACET-II.

Property	**Drive Beam**	**Witness Beam**
Charge	1.5 nC	0.5 nC
Energy	10 GeV	10 GeV
ϵ_{nx}	3.4 μm rad	3.2 μm rad
ϵ_{ny}	3.0 μm rad	3.1 μm rad
$\beta_{x,y}$	70 cm, 70 cm	5 cm, 5 cm
$\alpha_{x,y}$	4.2, 1.6	0, 0
σ_z	5.2 μm	5.2 μm
σ_δ	0.1%	0.1%
Δz	150 μm	

of the plasma lens, and ends 600 μm after the center of the lens. A timestep of 1.08×10^{-15} s is calculated from the Courant condition.

The waist location of the witness beam is set to be at the center of the plasma lens. The plasma density profile in the longitudinal z and vertical y dimensions is loaded as a 2D fit from a laser-ionized plasma density profile using techniques introduced in Chapter. 4, and the density in x assumed constant across the window dimensions. The 2D fit is a double-tanh profile that is elliptical in the $y-z$ plane and includes a slight linear gradient in y due to using a realistic gas jet density profile:

$$n(\vec{r}) = n_0 \left(\frac{1}{2} + \frac{1}{2}\tanh\frac{\zeta + a}{b}\right)\left(\frac{1}{2} - \frac{1}{2}\tanh\frac{\zeta - a}{b}\right)(dy + 1) \tag{3.45}$$

where

$$\zeta(y, z) = \sqrt{(cy)^2 + z^2}, \tag{3.46}$$

$n_0 = 2.9411{\times}10^{16}$ cm^{-3}, $a = 376.41$ μm, $b = 68.015$ μm, $c = 4.5828$, and $d = -4.5298{\times}10^{-4}$ μm^{-1}. Parameter a is effectively the HWHM of this density distribution, and was chosen to recreate the necessary plasma lens thickness used in the example of Sec. 3.4.1. We will discuss linear density gradients more in Chapter 4.2 when we go into more detail on how to generate plasma lenses, but in this particular simulation the off-axis variation was quite small in the transverse plane. Even while including this more realistic transverse density profile, the lens essentially behaves as if it had a uniform transverse density profile.

Particles within the beam are initially weighted from 0 to 1 in a Gaussian distribution. Once the witness beam has exited the plasma, we extract its phase space. The weighted electrons are then imported into our particle propagation code and the beam is evolved through vacuum to the waist position. This two-step procedure allows for VSim to capture the complex dynamics of the plasma interaction, while saving significant computation time for modeling vacuum propagation.

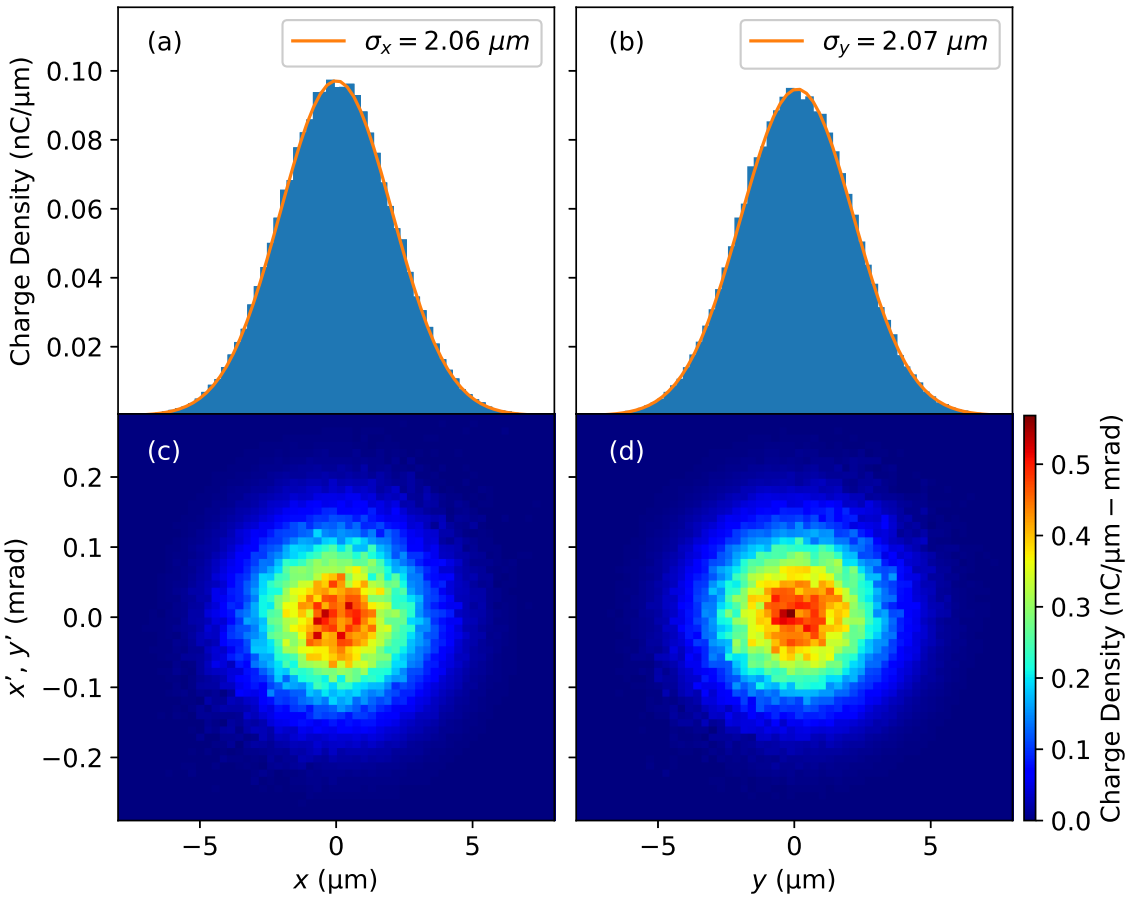

Figure 3.14: Projected distribution of the horizontal (a) and vertical (b) coordinates of the electron witness beam at its waist following the passive thin plasma lens. A Gaussian fit to the profile is shown with the orange line. Transverse phase space in the horizontal (c) and vertical (d) planes exhibit perfectly Gaussian profiles in all transverse dimensions.

Figure 3.14 shows the final transverse phase space distribution of the witness beam in the x and y planes at the waist following the plasma lens. The statistical rms size of the electron beam is 2.04 μm and 2.05 μm in the x and y dimensions, respectively. These values differ by less than

1% from the those of a Gaussian fit, 2.06 μm and 2.07 μm, as shown in Fig. 3.14. The vacuum waist size without the lens is 2.86 μm and 2.81 μm in x and y. The lens has therefore decreased the minimum spot size by a factor of $1/\sqrt{2}$, corresponding to the intended halving of the beta function. The final waist location obtained from the simulations is 2.47 cm downstream of the lens, which agrees well with the thin lens theoretical calculation of 2.49 cm. The final beam distribution in transverse phase space is symmetric and aberration-free, as demonstrated by the perfectly Gaussian distributions shown in Fig. 3.14. The witness beam's transverse emittance in the horizontal and vertical dimensions grew from 3.199 μm rad to 3.201 μm rad and from 3.100 μm rad to 3.104 μm rad, respectively. The change in the witness beam's energy and energy spread in the lens is negligibly small with centroid energy $\gamma_b = 19570.5$ and rms energy spread 0.10010% before the plasma lens, and centroid energy $\gamma_b = 19580.7$ and rms energy spread 0.10008% after the plasma lens.

This simulation demonstrates that under ideal situations the underdense plasma lens is a capable tool for focusing electron beams with strong, axisymmetric, and aberration-free focusing. The simulation results also agree with the analytic modeling presented in this chapter.

3.6.2 Second Wake Period Focusing

Moving on from ideal plasma lens operation, let us consider a special case that can allow the plasma lens to reach higher densities than would normally be possible. Up to this point, we have been assuming that the trailing witness electron bunch is fully within the first wake period behind the leading drive electron bunch. This includes the assumption that the longitudinal spacing between the witness and drive bunches is less than the longitudinal extent of the wake bubble, which roughly corresponds to the plasma wavelength. However, this bubble immediately behind the drive beam is not the only region where the plasma electrons are driven off-axis. In the highly nonlinear blowout regime, the wake structure itself is periodic so there is a train of bubbles behind the first blowout wake.

It is possible to consider a scenario where the spacing between the drive and witness bunches are such that the witness bunch is placed within a later wake bubble. This would allow the plasma

lens to operate at a higher density than would be typical given the constraints between longitudinal separation and plasma density. The only further requirement left is that the witness bunch is small enough in size to fit within the second wake bubble, which is generally smaller than the first bubble.

This regime of second bubble focusing is far from ideal, however. Wake bubbles behind the first are much harder to predict analytically, as the plasma dynamics at the rear of the first wake are highly nonlinear and can include complex processes where the plasma density is large and velocities are high. This leads to the second and later wake bubbles having less sharp sheath boundaries and increases the possibilities of aberrations being included. Furthermore, if one imagines a density transition for such a plasma lens between vacuum and the high density region of second wake bubble focusing, then there will be a density transition where the rear of the first wake bubble overlaps with the witness bunch at some longitudinal position within the plasma lens. This overlap would be unavoidable and most certainly induce aberrations in the focusing.

With these pros and cons in mind, we can investigate this interesting regime with a proof-of-principle PIC simulation. To reflect a typical situation where one would perhaps want to use a very thin but very dense plasma lens, we choose a plasma lens with a thickness of 72 μm and peak density 3.7×10^{17} cm^{-3}. The drive-witness separation is 100 μm, and other beam parameters are chosen as in the previous section. The result is a plasma lens that operates with both the plasma lens thickness and longitudinal wake length that are less than the drive-witness separation distance, as seen in Fig. 3.15.

With the same two-step methodology from the previous section, the 6D phase space of the witness bunch's electrons are exported from the PIC simulation after the plasma lens and imported into a much faster particle tracking code. The witness bunch is then propagated through vacuum to its focus (Fig. 3.16), and at the focus we can plot the transverse phase space and see the witness bunch quality (Fig. 3.17). In this particular case with the correct plasma density, the focus quality is close to an aberration-free Gaussian. The emittance growth is also small, with the normalized emittance in only increasing from 3.15 to 3.20 mm $-$ mrad. The slight asymmetric focusing seen in Fig. 3.16 is due to the slight asymmetric initial emittances for the witness bunch, as written in

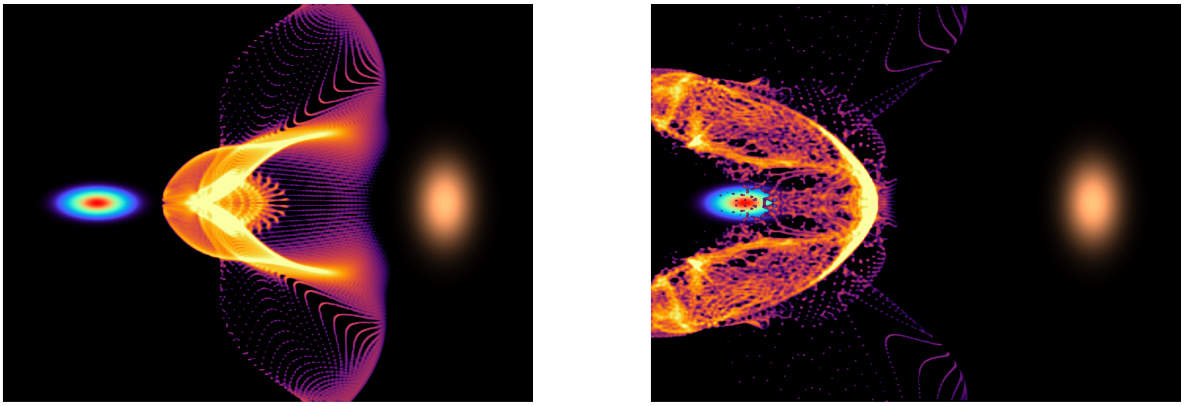

Figure 3.15: Two PIC snapshots of a plasma lens operating with the witness bunch in the second wake bubble. Both images show the x-z plane along the $y = 0$ axis. The relative plasma density from low (purple) to high (yellow), the Gaussian drive bunch is shown in orange, and he Gaussian witness bunch is shown in rainbow blue-to-red. Both beams are propagating towards the right. The left image shows the 40 μm plasma after the drive bunch has passed and before the witness bunch arrives, and the rear of the first wake is located between the two bunches. The right image is slightly later in time, where the plasma has evolved and the witness bunch is within the second wake bubble.

Table 3.3.

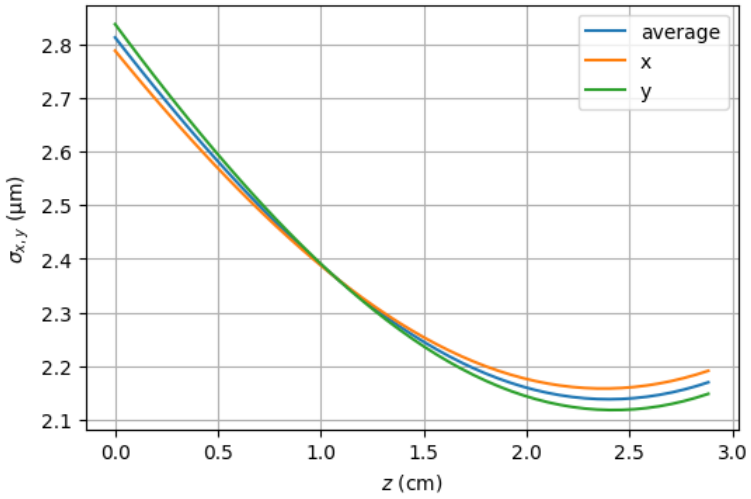

Figure 3.16: Beam spot size evolution after the second wake bubble plasma lens. Both transverse rms spot sizes are given for x (orange) and y (green), as well as the average of the two (blue). The slight asymmetric focus is due to the initial asymmetry in the electron beam.

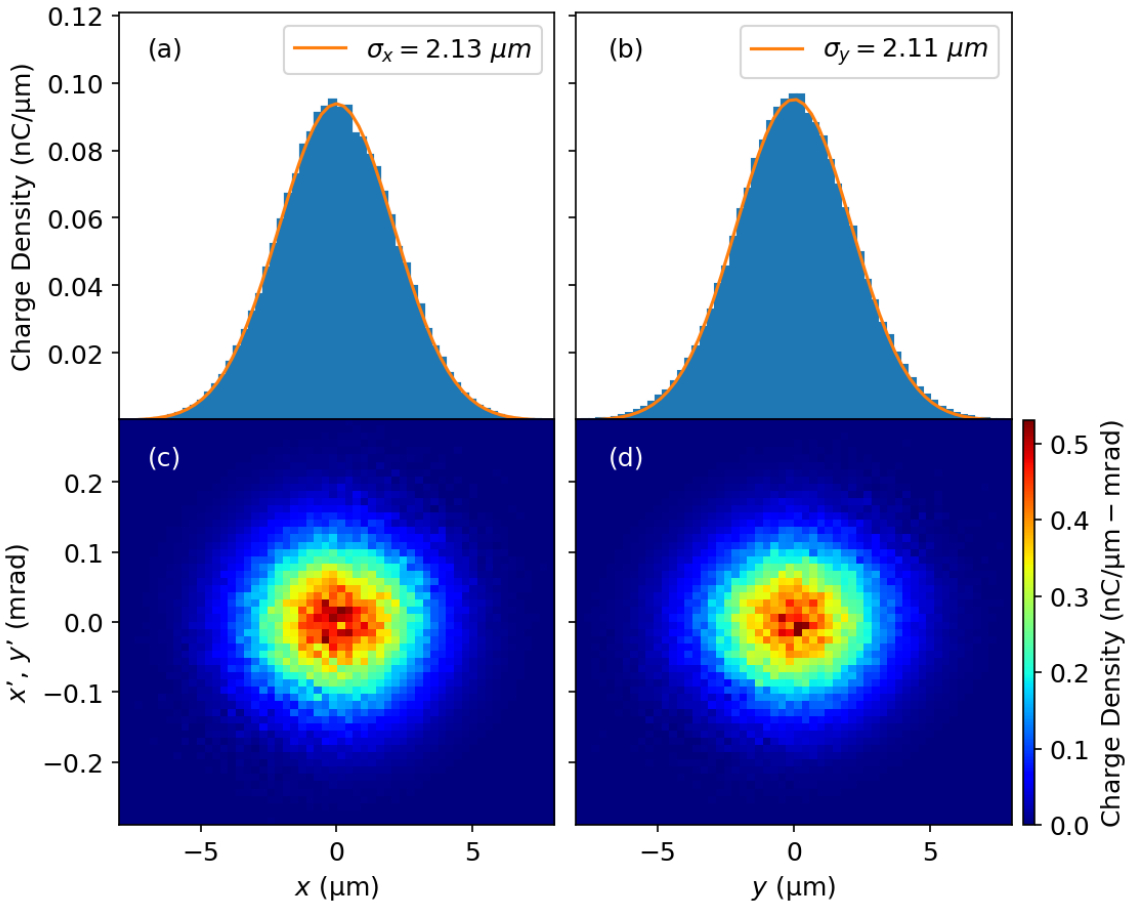

Figure 3.17: Transverse phase space of the electron beam at the focus following a second wake bubble plasma lens. The beam distributions follow a Gaussian profile with little aberrations present.

Part of the reason this particular simulation works so well is that the initial plasma density profile is quite sharp. The density ramp transition between vacuum and peak density is small, so there is only a small longitudinal region of a transitional plasma density where a wake bubble forms with its rear at the witness bunch location. This is never an issue in the ideal plasma lens operation because the witness bunch is located within the first bubble. Any plasma density lower than this will only increase the bubble's size and the electron bunch will still be within the first wake period.

The peak plasma lens density itself is also quite critical and dependent on the drive-witness separation. For a separation of 100 μm, a plasma density of 3.7×10^{17} cm^{-3} is ideal. At a slightly lower plasma density of 2.0×10^{17} cm^{-3}, the first bubble closes right on the witness bunch, and the

high local plasma electron density at this location subsequently blows apart the witness bunch. On the other hand, one could consider increasing the plasma density until the witness bunch is within the third bubble. However, in the parameters chosen here the plasma bubble becomes too small for the witness bunch to fit and the plasma density becomes too large for the drive bunch to drive a strong blowout wake.

Overall, this brief simulation study has shown that with the correct parameters a plasma lens operating in the second wake bubble can be viable. It is an attractive method for operating plasma lenses at densities greater than what would be possible with limited electron beam parameters, and could be a useful regime for particularly thin plasma density profiles.

3.6.3 Single-Bunch Plasma Lens Focusing

Moving on from plasma lens operation using a separate drive and witness bunch, now we consider the case where we are operating with only a single electron bunch. Rather than using two electron bunches where the first bunch drives the nonlinear blowout wake and the second bunch witnesses the strong focusing forces, here we have a single long electron bunch. Given a large enough beam density, the head of this single bunch can drive the nonlinear blowout wake and the rest of the beam will receive the focusing from the ion column.

As a proof-of-principle, we perform PIC simulations to demonstrate the process of the single bunch underdense plasma lens regime and evaluate the electron beam's quality after the plasma lens. Electron beam parameters are chosen similar to the drive beam parameters in Table 3.3, although with no witness bunch and a slightly longer drive bunch. We choose here a plasma density of 10^{17} cm^{-3} and thickness of 50 μm, and the resulting blowout wake around the electron bunch is shown in Fig. 3.18.

Using these parameters, the head of the electron bunch is capable of driving a sufficiently wide blowout wake in the plasma. The electron bunch itself is small enough to be contained within the blowout, and so a large fraction of the electron bunch will witness the high focusing forces without strong aberrations. For this particular simulation, the resulting transverse beam distribution at the

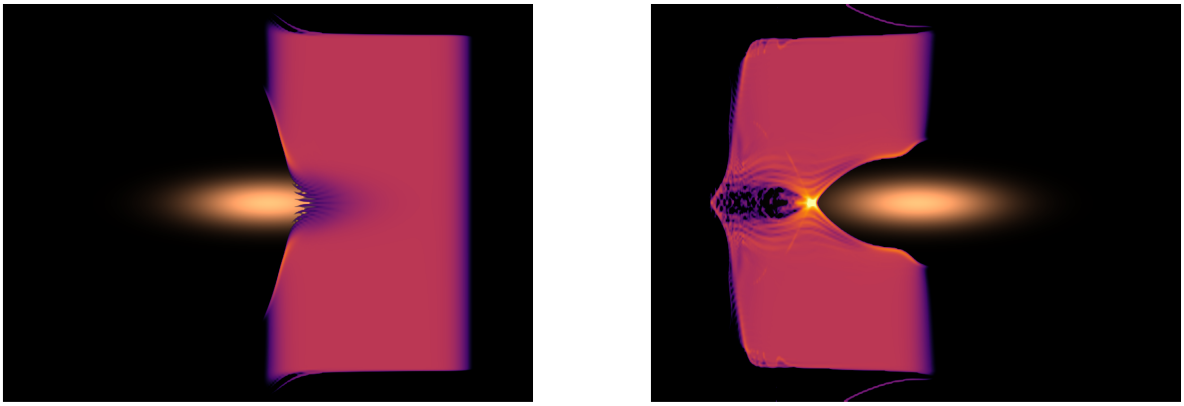

Figure 3.18: Two PIC snapshots of a plasma lens operating with a single electron bunch. Both images show the x-z plane along the $y = 0$ axis. The relative plasma density from low (purple) to high (yellow), the Gaussian electron bunch is shown in orange. The electron beam is propagating towards the right. The left image shows the 50 μm plasma right as the head of the electron bunch is driving a blowout wake. The right image is slightly later in time, where the plasma has evolved and the bulk of the electron bunch has been within the blowout wake.

subsequent focus is plotted in Fig. 3.19. While the profile for the full beam is no longer Gaussian, the distribution can instead be approximated as the sum of two Gaussian functions. The inner core makes about 70% of the electrons for this particular case and represents the electrons in the bunch that were strongly focused by the ion column in the blowout wake. The other 30% of the electrons are included in the larger, outer Gaussian. These electrons were either at the head of the bunch and drove the blowout wake, or they were overlapping with the blowout sheath and received focusing aberrations.

While this is not an ideal scenario for the plasma lens, it is nonetheless much easier to realize in experiment simply due to requiring only a single electron beam. This single bunch regime allows for plasma lenses to be tested while an accelerator is still building up its capabilities for two electron bunches, for example. However, this regime is not without its downsides. Since the single electron bunch is both driving the wake and being focused, the beam as a whole has many focusing aberrations induced during the interaction and it is difficult to discern the head and tail of the beam.

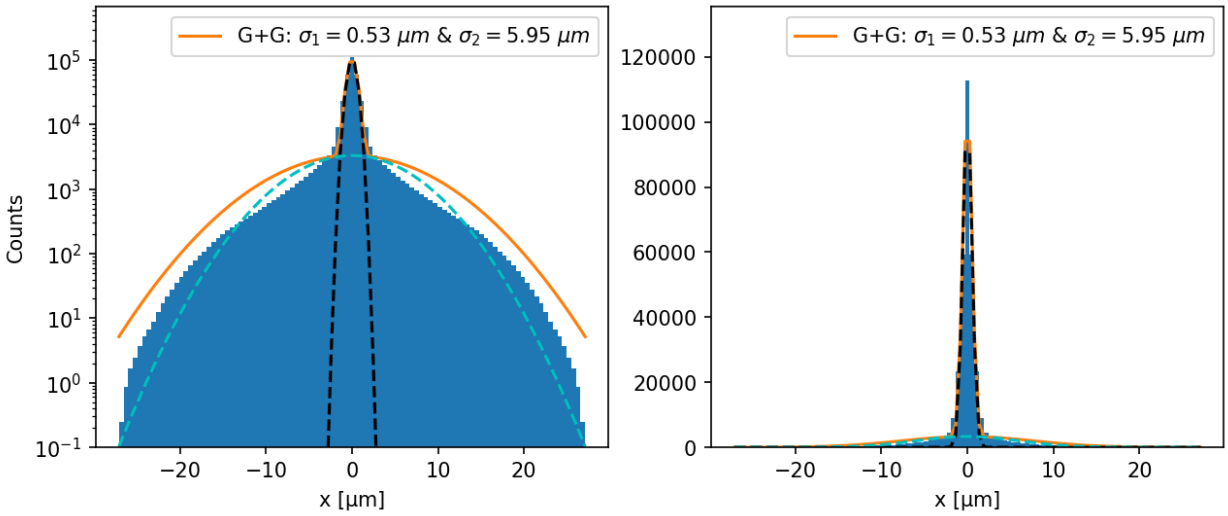

Figure 3.19: Transverse beam distribution at the focus of a plasma lens operating in the single bunch regime, plotted in a logarithmic scale (left) and linear scale (right). The beam's distribution can be approximated as the sum of two Gaussian functions (solid orange), which includes a large rms spot size from electrons not focused (dashed blue) and a tight core of focused electrons with a small rms spot size (dashed black).

Chapter 4

Passive Plasma Lens Design

Now that we have a general idea for how a plasma lens operates and what parameters the focusing force and aberrations scale with, we can begin to examine methods to generate such a plasma lens. In particular, we are interested in generating a plasma that has a density on the order of $n_p \sim 10^{16} - 10^{19}$ cm^3 with a thickness of the order of $L \sim 10's - 100's$ μm. Furthermore, to serve as an application for focusing electron beams this plasma lens must exist in an electron beamline. The obvious requirement here is that there needs to be a trajectory through the plasma lens for the electron beam to pass through, though there may be other requirements which have to be addressed: such as the requirements on the vacuum level, repetition rate of the incoming electron beams, and tolerance of equipment to radiation levels. In this section, we analyze the scheme that appears to be the most useful for applications at FACET-II: using a laser to ionize the outflow of a gas jet.

First we discuss the ionization process, and offer schemes to use a laser to ionize a plasma lens with appropriate dimensions. Next, we discuss how a gas jet outflow's density profile impacts the generation of a plasma lens. Lastly, we analyze how the transverse density nonuniformity present in a gas jet outflow can impact the beam dynamics when used as a plasma lens.

4.1 Plasma Generation

Here we introduce two methods of ionizing a neutral gas into a plasma: laser ionization and electron beam ionization. Then, for laser ionization, we discuss two schemes of ionizing a plasma

lens. The first is that of a spherical lens focus, which is easier to implement and will be tested first at FACET-II. The second is at the focus of crossed cylindrical lenses, which are more difficult to implement but will result in a more favorable density profile. We then discuss how the ionizing laser refracts as it ionizes the plasma and how this impacts the plasma lens generation.

4.1.1 Laser Ionization

Neutral atoms can be ionized by a laser to form a plasma. The ionization process is dictated by the ionization energy of the atoms and the intensity of the laser's electric fields. This process is described in three different regimes depending on the relative strength of the laser's electric field. From low to high intensity, these regimes are the multi-photon, tunneling, and barrier suppression ionization (BSI) regime (Fig. 4.1).

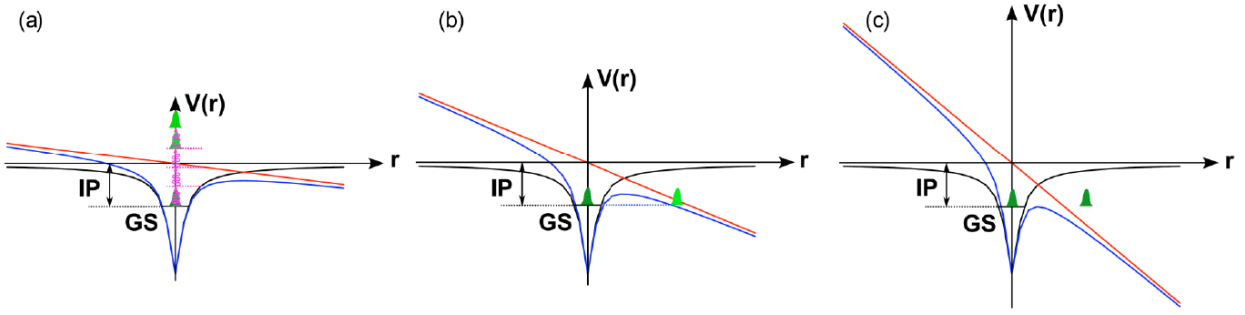

Figure 4.1: Ionization regimes as determined by the applied electric field strength. (a) multi-photon regime: a single photon does not have enough energy so multiple photons collide with a bound electron to ionize. (b) tunneling regime: applied field lowers the potential well so that there is a nonzero probability for the electron to tunnel out. (c) barrier suppression regime: the electric field distorts the potential well to completely unbound the electron. Reproduced from Ref. [73].

In the multi-photon ionization regime, the electric field from the laser is not intense enough to allow for tunneling of an atom's electrons out of the bound state but the electrons can still absorb the energy of several photons to ionize. If we look at an engineering formula for the energy of a single photon,

$$E(eV) = \frac{1.2398}{\lambda(\mu m)}, \tag{4.1}$$

we find that for a typical Ti:Saphh laser system operating at 800 nm the energy of a single photon

is about 1.55 eV. We can compare this to the ionization energy for hydrogen gas, H_2, at 15.4 eV, which reveals that ten photons would need to be simultaneously absorbed to singly-ionize one molecule of H_2. The ionization rate and resulting ionization fraction of the full laser pulse interacting with a region of neutral gas is described in the PPT ionization theory for this multi-photon ionization regime [74].

In the tunneling regime, the electric field from the laser significantly distorts the atomic potential of the neutral atom or molecule. This allows for a bound electron to have a chance at tunneling out of the atomic potential well and thus ionizing into a plasma. As the electric field becomes stronger, the probability for this tunneling event increases. To calculate the full ionization rate for the interaction with laser, one can either use the PPT theory or the ADK theory[75]. The ADK ionization theory is an approximation on the PPT formalism that is only valid in the tunneling regime, but it is a more straightforward calculation. Tunneling ionization is the primary regime we will be focusing on in this chapter.

Lastly, in the BSI regime the electric fields from the laser pulse are intense enough as to completely remove the barrier for a bound electron to escape its atom or molecule. This is qualitatively as if the thickness of the barrier in the tunneling regime has gone to zero. However, special care must still be taken to accurately predict the ionization rate in this regime. Both the PPT and ADK ionization models are not valid above the BSI threshold, and instead one must use time dependent Schrodinger equation (TDSE) simulations to numerically calculate the ionization rate [76]. For the remainder of this work, we will be operating in regimes under the BSI threshold where the intensity is low enough to remain in the tunneling regime.

For the purposes of the plasma lens ionization discussed in this chapter, we will primarily be interested in the ADK model of ionization. While the ADK model is only valid in the tunneling regime, this is the only regime that we are concerned about. The low density contributions that arise from multi-photon ionization are irrelevant for the short, sharp plasma density profiles of plasma lenses, and so the more straightforward ADK model is appropriate.

We start with the ionization rate in the ADK model for a constant electric field. From

Ref. [74], this is given as

$$\omega_{stat}(E) = |C_{n^*l^*}|^2 \frac{(2l+1)(l+|m|)!}{2^{|m|}(|m|)!(l-|m|)!}\xi_i \left(\frac{2E_0}{E}\right)^{2n^*-|m|-1} \exp\left(-\frac{2E_0}{3E}\right). \tag{4.2}$$

Equation 4.2 is written in atomic units, where $m_e = \hbar = e = 1$. The ionization energy for a valence electron is ξ_i, and its angular momentum and projection are l and m, respectively. The intra-atomic electric field E_0 is calculated from the electrom momentum $\kappa = \sqrt{2\xi_i}$ such that $E_0 = \kappa^3$. The external electric field E is assumed small compared to the intra-atomic electric field. $n^* = Z/\kappa$ is the effective principle quantum number, Z is the charge of the atomic residue which, for ionization of a neutral atom, will be $Z = 1$. The constant $C_{n^*l^*}$ is calculated in the Coulomb wave function using effective quantum numbers:

$$|C_{n^*l^*}|^2 = \frac{2^{2n^*}}{n^*\Gamma(n^*+l^*+1)\Gamma(n^*-l^*)}. \tag{4.3}$$

In Eqn. 4.3, $l^* = n_0^* - 1$ with the effective principal quantum number of the ground state $n_0^* = n^*$ since we are not interested in ionization from excited states[77, 78].

Interestingly, the ADK theory as written uses an approximation for Eqn. 4.3 that is written as

$$|C_{n^*l^*}|^2 \approx \left(\frac{4e}{n^{*2}-l^{*2}}\right)^{n^*} \left(\frac{n^*+l^*}{n^*-l^*}\right)^{l^*+1/2} \frac{1}{2\pi n^*}, \tag{4.4}$$

but this approximation has significant errors so it is not used in practice [79].

While Eqn. 4.2 is valid for a constant electric field, we now need to consider the ionization rate for a laser pulse's time alternating electric field:

$$E(t) = E\cos(\omega t) \tag{4.5}$$

with laser frequency ω. This frequency must be small when compared with the tunneling frequency in order to assume that the electric field is constant during the time it takes for an electron to tunnel through the potential barrier. Integrating over Eqn. 4.5 for linearly polarized light gives an added constant for the ionization rate:

$$\omega_{lin}(E) = \left(\frac{3E}{\pi E_0}\right)^{1/2} \omega_{stat}(E). \tag{4.6}$$

Lastly, since the ionization rate for the $m = 0$ state is larger than for other states by a factor of

$$\frac{\omega_{m=0}}{\omega_{|m|=1}} = \frac{2E_0}{E} \tag{4.7}$$

and we already assumed the intra-atomic electric field E_0 is much larger than the laser's electric field E, we only need to consider the $m = 0$ states for calculating the ionization rate.

To make Eqn. 4.2 more approachable, we convert the atomic units back to practical units of SI. The one exception is the ionization energy ξ_i, which we will write in terms of eV. The ionization energy in Hartree is replaced with

$$\xi_i[h] \rightarrow \frac{e}{\alpha^2 c^2 m_e} \xi_i[eV], \tag{4.8}$$

ionization rate is replaced with

$$\omega[1/\tau_a] \rightarrow \frac{\hbar}{\alpha^2 c^2 m_e} \omega[1/s], \tag{4.9}$$

and field strength is replaced with

$$E[E_a] \rightarrow \frac{e\hbar}{\alpha^3 c^3 m_e^2} E[V/m]. \tag{4.10}$$

Field E_0 is changed to be given in V/m with

$$E_0[V/m] = 2^{3/2} \frac{\sqrt{m_e e}}{\hbar} \xi_i^{3/2}[eV] \tag{4.11}$$

and n^* has an additional factor arising from $\xi_i[eV]$:

$$n^* = \alpha c \sqrt{\frac{m_e}{2e}} \frac{Z}{\sqrt{\xi_i[eV]}}. \tag{4.12}$$

Using these substitutions, the ionization rate for the ADK model of Eqn. 4.2 can instead be written as

$$\omega_{stat}(E) = |C_{n^* l^*}|^2 \frac{e}{\hbar} \frac{(2l+1)(l+|m|)!}{2^{|m|}(|m|)!(l-|m|)!} \xi_i \left(\frac{2E_0}{E}\right)^{2n^*-|m|-1} \exp\left(-\frac{2E_0}{3E}\right). \tag{4.13}$$

Example calculations using Eqn. 4.13 to calculate the ionization fraction are shown in Fig. 4.2 for various neutral gases. For the remainder of this chapter, we will be assuming diatomic Hydrogen H_2 as the gas of choice for the plasma lens, as the the low ionization energy makes it easier to ionize with a small budget of laser energy.

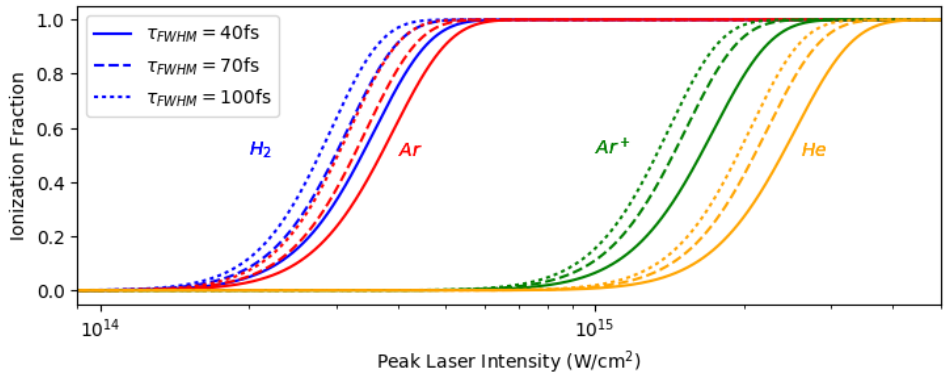

Figure 4.2: Ionization fraction for various several gas species as a function of peak laser intensity for a Gaussian beam. The colors represent the gas species and the line styles represent different temporal Gaussian pulse widths, as given by the legend.

4.1.2 Beam Ionization

In addition to laser ionization, a neutral gas can also be ionized by the intense electric fields of the electron beam itself. This, too, can be modeled using ADK ionization for the tunneling regime where the electric field is instead given by the electron beam's Coulomb field. We first consider the limit where the electron beam is relativistic and we are calculating the electric field at a far enough radial distance r such that $\sigma_r \ll r \ll \gamma\sigma_z$, where σ_r and σ_z are the electron beam's rms bunch radius and length, respectively. In this limit, the electron beam acts as a line charge and the electric field is simply

$$E_r = \frac{\lambda(z,t)}{2\pi\epsilon_0 r}, \tag{4.14}$$

with λ as the linear charge density.

To calculate the ionization rate, we use Eqn. 4.14 in conjunction with the ionization rate for a static electric field (Eqn. 4.13). As this Coulomb field is not oscillating, there is no need for the additional term from Eqn. 4.6. For particular beam parameters, this can even be as large as the electric fields in a laser pulse designed for ionization. Of course, this only applies for when $r \gg \sigma_r$. Even if we assume that the Coulomb field peaks when $r = \sigma_r$, within the electron beam we would also need to consider ionization through collisions between the gas and electron beam. Examples of

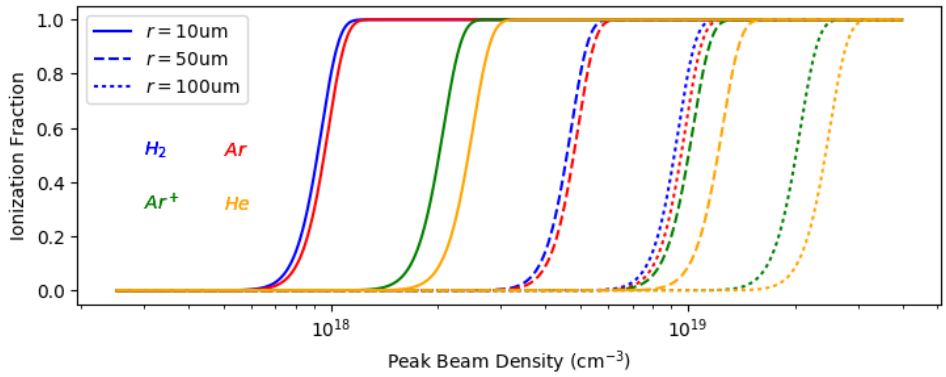

Figure 4.3: Ionization fraction for various several gas species as a function of peak number density for an electron beam with a Gaussian density profile. The colors represent the gas species and the line styles represent several distances from the electron beam axis, as given by the legend. The transverse beam size is assumed $\sigma_{x,y} = 5$ μm and the longitudinal size is assumed $\sigma_{x,y} = 20$ μm.

using Eqn. 4.14 to calculate the ionization rate for Gaussian electron beams are shown in Fig. 4.3.

Beam ionization can either be harmful to a plasma lens's design, or alternatively one can specifically design a plasma lens where the plasma is deliberately designed to be beam ionized. With regards to the former, beam ionization can significantly increase the thickness of a plasma lens if there is neutral gas beyond the desired thickness of the plasma lens. This is especially true if one would like to make a plasma lens that is a few $10's$ μm thick, as there is no known way to contain a neutral gas in such a small area open to vacuum on all sides. For example, if we are using a laser to ionize a 100 μm thick plasma lens within the outflow of a gas jet nozzle there is still plenty of unionized neutral gas both upstream and downstream of the pre-ionized plasma. To mitigate this, one would have to look at the parameters of the driving electron beam and pick a gas species for the gas jet that has a ionization energy larger than the electron beam but smaller than the laser's intensity.

On the other hand, beam ionization can be a useful tool in making a plasma lens. From Sec. 3.2, we found that for a thin lens the focal length is calculated from the integrated plasma density. Thus, we can have a plasma lens that spans the entire width of a gas jet and, if the density is low enough, the plasma would still be considered a thin lens. One can then choose parameters so

that a drive beam would be able to ionize the neutral gas and drive a blowout wake for the width of the full radial extent of the gas jet outflow. The downside to such a design is that we take away all of the control a laser-ionized design would give, and instead the gas density profile itself would determine the plasma lens parameters.

4.1.3 Spherical Lens Focus

Now that we have a description for the ionization profile given a volume of neutral gas and a laser, we move to considering how to shape the laser pulse in order to ionize a plasma lens of the desired dimensions. The main requirements we want to satisfy is that (1) plasma density is as uniform as possible, (2) the longitudinal thickness on the order of $10's - 100's$ μm, and (3) the transverse plasma density profile is wider than the blowout wake in the plasma. Requirement (1) is an effort to replicate the ideal plasma lens operating conditions, and effectively means that we want the bulk of the plasma lens to be fully ionized. Requirement (2) determines the focusing strength of the plasma lens, and for this we want a scheme that focuses the laser accurately and is stable over many shots. Requirement (3) is, in addition to (1), another condition to be in the ideal blowout regime. If the blowout wake is wider than the plasma, then the blowout wake will become elliptical rather than circular and the focusing in the transverse planes will no longer be axisymmetric.

There are countless ways to focus a laser pulse to be able to achieve these specifications, but in this chapter we will consider two simple cases that are sufficient. The first is using a simple spherical lens to focus an incoming laser pulse to a Gaussian spot size. This spherical lens can either be a typical transmissive lens or a reflective off-axis parabolic mirror (OAP). Transmissive lenses are cheaper and easier to align than OAP's, but one needs to account for the added distortions when a high intensity laser pulse propagates through a transmissive material. A transmissive optic also has a wavelength-dependent focus so the focus of a transmissive optic will be in a finite region, but a reflective OAP focuses all wavelengths of the laser pulse to the same point for a higher peak intensity spot. However, with the goal of plasma lens ionization we only need to have a large

enough intensity in the focus so these laser focusing aberrations can be low enough to not impact ionization of the plasma lens.

To model the propagation of a Gaussian laser pulse through optics and free space, we use the complex beam parameter q defined in Ref. [80] as

$$\frac{1}{q(x)} = \frac{1}{R(x)} - \frac{i\lambda_0}{\pi n \omega(x)^2}, \tag{4.15}$$

where R is the phase front's radius of curvature, ω is the transverse spot size, λ_0 is the vacuum wavelength, and n is the index of refraction. The complex beam parameter can be solved for along the laser propagation axis x by evolving the complex beam parameter according to transfer matrix formalism. For a given transfer matrix of an optical system

$$M = \begin{bmatrix} A & B \\ C & D \end{bmatrix}, \tag{4.16}$$

the complex beam parameter can be evolved from q_i to q_f using

$$\frac{1}{q_f} = \frac{C + D/q_i}{A + B/q_i}. \tag{4.17}$$

Equation 4.16 can thin and thick lenses, flat and curved interfaces, and drift spaces.

So long as we can describe the transfer matrix of the optical system and the initial Gaussian parameters, we can also know the transverse spot size of the Gaussian laser pulse throughout the optical system. From this, around the focus for laser propagation axis x and radial distance $r^2 = y^2 + z^2$ we can determine the electric field using the fundamental transverse Gaussian mode (Ref. [81])

$$\mathbf{E}(r,x) = E_0 \hat{\mathbf{y}} \frac{\omega_0}{\omega(x)} \exp\left(\frac{-r^2}{\omega(x)^2}\right) \exp\left(-i\left(kx + k\frac{r^2}{2R(x)} - \psi(x)\right)\right) \tag{4.18}$$

and calculate the intensity of the laser pulse:

$$I(r,x) = I_0 \left(\frac{\omega_0}{\omega(x)}\right)^2 \exp\left(\frac{-2r^2}{\omega(x)^2}\right). \tag{4.19}$$

Above, $k = 2\pi n/\lambda$ is the laser wavenumber and $\psi(x) = \arctan(x/z_R)$ is the Gouy phase, with Rayleigh length $z_R = \pi\omega_0^2 n/\lambda$.

While the complex beam parameter is a quick and effective method to model the propagation of a laser pulse, it is only truly applicable for ideal Gaussian beams. This is not always the case, and many experiments using a Ti:Saphh laser system the laser pulse's intensity profile is typically closer to a flattop profile than a Gaussian. In an effort to make this complex beam parameter more general, we include the beam quality factor M^2 in the model to effectively increase the spot size.

As an example, Fig. 4.4 shows the propagation of a Gaussian laser pulse's spot size after a spherical lens or OAP. The initial laser parameters are chosen to be similar to the laser system of FACET-II, which will be discussed in the following chapter. The laser wavelength is $\lambda = 800$ nm, with energy $E = 10$ mJ and pulse duration $\Delta\tau = 70$ fs. The initial Gaussian spot size is 13.6 mm, which is roughly equivalent to a beam with 30 mm flattop intensity profile. A beam quality factor of $M^2 = 2$ is chosen to roughly estimate the deviation of a realistic laser from a true Gaussian profile.

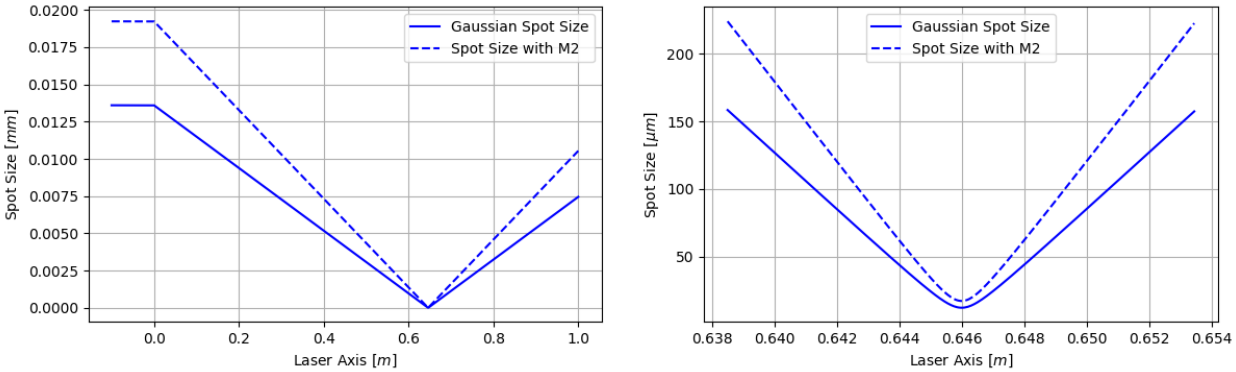

Figure 4.4: Propagation of a Gaussian laser spot size (solid blue) of wavelength $\lambda = 800$ nm after a spherical lens of $f = 646$ mm located at $z = 0$. The spot size at focus is 12.1 μm. Plotted in dashed blue is the inclusion of a beam quality factor of $M^2 = 2$, which effectively increases the spot size by a factor of $\sqrt{M^2}$ while keeping the Rayleigh length the same.

The laser pulse is propagated through a spherical lens of $f = 646$ mm located at $z = 0$. This can either be a typical thin lens or an OAP, which would have additional benefits like focusing all wavelengths to the same point and not distorting the laser by propagating through the lens itself. This optic focuses the Gaussian laser pulse to a spot size of 12.1 μm, and including the M^2 factor this spot size is approximately 17.1 μm. The intensity achieved at the focus is more than enough

to ionize diatomic Hydrogen gas (H_2) with an ionization energy of 15.4 eV, so much so that under this ideal situation the H_2 is ionized upstream and downstream of the focus where the spot size is larger. The ionization profile is plotted in Fig. 4.5. Here, it is intended for an electron beam to travel on an axis orthogonal to the laser's propagation axis. The laser travels in the x direction, which is horizontal to the electron beam's axis in the z direction.

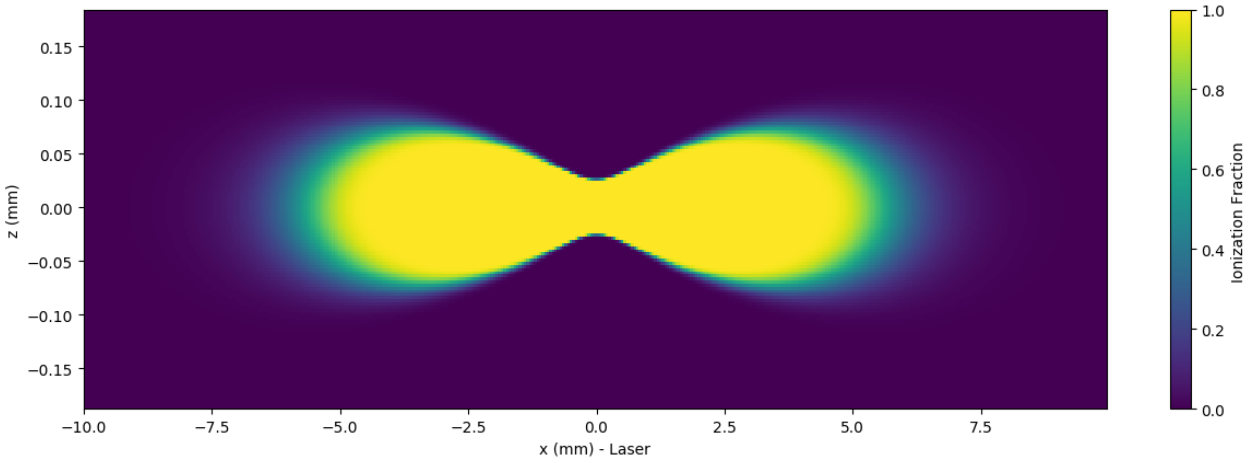

Figure 4.5: Ionization profile around the laser focus of a spherical lens. The x axis is the laser propagation axis and the z axis is a transverse axis, although the intensity profile here is axisymmetric around x. In the color scale, yellow is fully ionized and purple is not ionized.

A benefit of the axisymmetric ionization profile of Fig. 4.5 for the purposes of plasma lenses is that there is a range of plasma lens thickness depending on where the electron beamline is located. If the laser's propagation is set up such that the focus is located on the electron beamline, then the plasma lens thickness is at a minimum of 53.9 μm. However, if the laser is instead set up such that the electron beamline is 4.05 mm upstream of the laser focus, then the resulting plasma lens thickness is much larger at 173.1 μm. The downside is that the longitudinal thickness and the transverse width of the plasma lens are the same, so an appropriate plasma density is needed for the blowout wake diameter to not exceed the width of the plasma lens.

Overall, this scheme of plasma lens ionization through a spherical lens focus is quite versatile despite the drawbacks of a smaller transverse width at the ionization profile's waist. The ability to access a wide range of plasma lens thicknesses makes this setup attractive as a first plasma lens

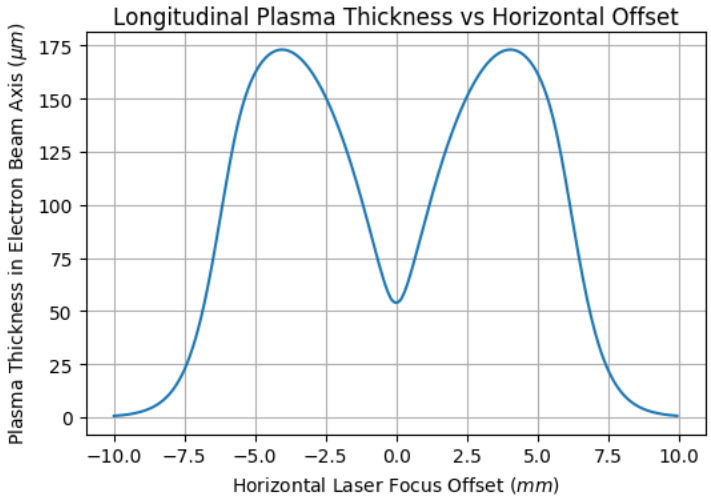

Figure 4.6: Range of plasma lens thicknesses vs the position along the laser's propagation axis for the plasma ionized at the focus of a spherical lens. The thickness of the plasma lens can be effectively adjusted by shifting the offset of the laser's focus with respect to the trajectory of the electron beam. In this example, a minimum plasma lens thickness of 55 μm is obtained when the laser focus and electron beamline are aligned, and a maximum of 175 μm is obtained at an offset of 4 mm.

experiment, and only requiring a single focusing optic makes the system easy to align. Refraction effects are negligible if used in conjunction with a gas jet, as we will see in analyzing refraction for the second example in Sec. 4.1.6 and Sec. 4.2.3. For these reasons, this spherical lens setup is what we have chosen to pursue at FACET-II, which will be discussed in Chapter 5.

4.1.4 Crossed Cylindrical Lenses

The second example of laser ionization we consider is using crossed cylindrical lenses instead of a single spherical lens. This gives us better control of the spot size, where now we can have one transverse spot size that is wide enough for the blowout wake and the other transverse spot size tailored to ionize a plasma lens of a particular thickness. For a laser pulse propagating in x, the transverse components z (horizontal) and z (vertical) are separable and can be modeled independently using their own complex beam parameter. The Gaussian electric field and intensity in Eqns. 4.18 and 4.19 are modified by letting $r^2 \to y^2 + z^2$ and $\omega(x)^2 \to \omega_y(x)\omega_z(x)$. Otherwise, the modeling of the laser pulse's propagation is the same as in Sec. 4.1.3.

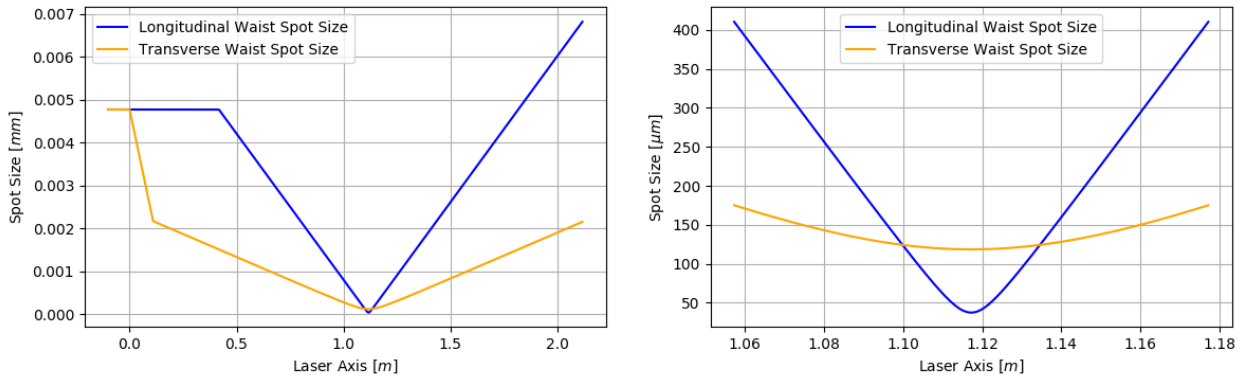

Figure 4.7: Propagation of a $\lambda = 800$ nm Gaussian laser spot size in the transverse horizontal (solid blue) and transverse vertical (solid orange) planes. The vertical laser component is focused by two cylindrical lenses of focal lengths $f_1 = 200$ mm and $f_2 = -100$ mm located at $x_1 = 0$ and $x_2 = 109$ mm, respectively. The horizontal laser component is focused by a single cylindrical lens of $f_3 = 700$ mm at $x_3 = 417$ mm. The smaller, 37.4 μm waist of the blue curve contributes to the longitudinal thickness of the plasma lens, while the larger, 118.5 μm waist of the orange curve keeps the plasma lens transverse width wide enough for the blowout wake.

A Gaussian laser propagating through three crossed cylindrical lenses is plotted in Fig. 4.7. The initial laser parameters are mostly the same as in Sec. 4.1.3, with two exceptions. The initial spot size is instead chosen to be 4.77 mm to correspond to a 10 mm flattop intensity profile, and the beam quality factor here is assumed to be $M^2 = 1$ for simplicity. The smaller initial spot size leads to the spot size at the focus being larger, but attention to the intensity at the focusing optics is needed to make sure the optics do not burn.

The crossed cylindrical lenses are set up so that the horizontal component is focused by two cylindrical lenses of $f_1 = 200$ mm and $f_2 = -100$ mm located at $x_1 = 0$ and $x_2 = 109$ mm, respectively. The vertical component is focused by a single cylindrical lens of $f_3 = 700$ mm at $x_3 = 417$ mm. Both components reach a focus at the same position, where the focus is 37.4×118.5 μm. Once again, if we consider diatomic Hydrogen gas H_2 the intensity profile is large enough to give a wide region where the gas should be fully ionized. The intensity profile along the central axis in all three planes is shown in Fig. 4.8.

Contrary to the spherical lens example of Sec. 4.1.3, in this setup the plasma lens is designed to be operated with the electron beam axis through the laser waist position. Here, the narrow waist

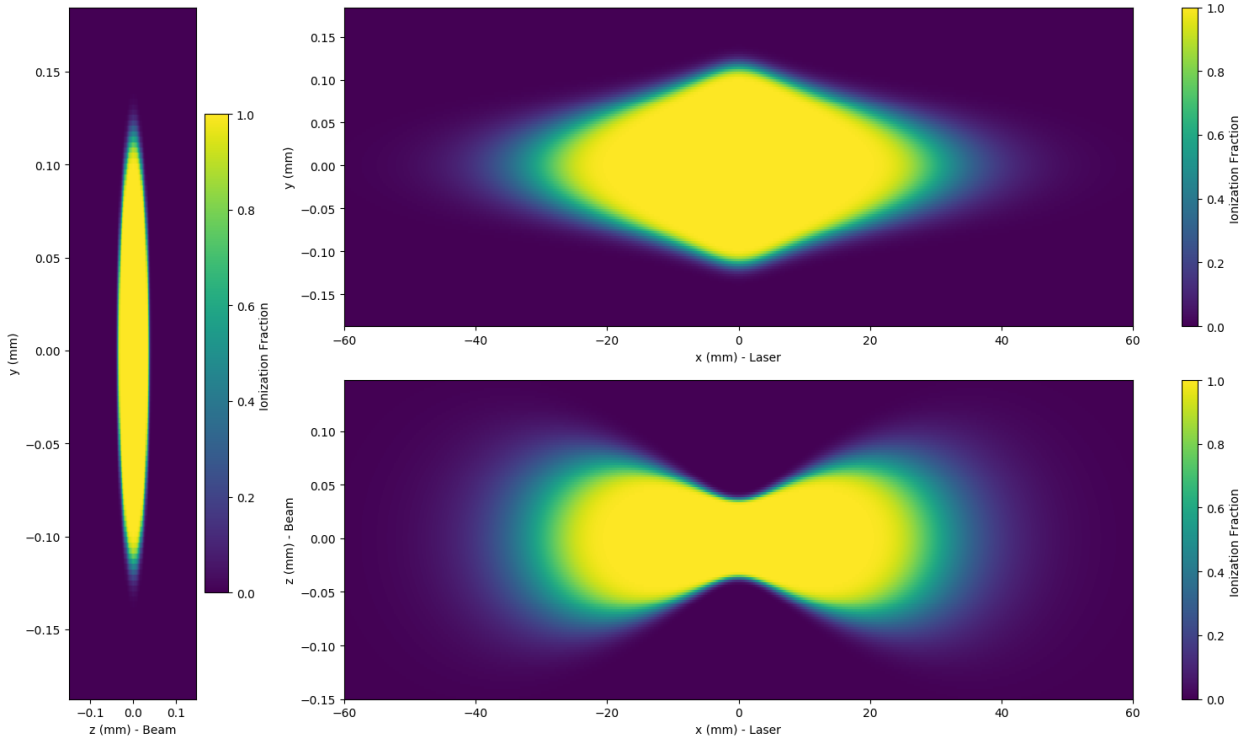

Figure 4.8: Ionization profile around the laser focus of crossed cylindrical lenses for the central axis of all three planes. In the color scale, yellow is fully ionized and purple is not ionized.

in the horizontal axis corresponds to the longitudinal thickness of the plasma lens, while the wider waist of the vertical axis is chosen to satisfy the condition that the plasma lens is wide enough to contain the blowout wake. In this example, the transverse plasma width is 200 μm while the longitudinal plasma thickness is 78.8 μm.

This optical setup can be more difficult to install and align given the three focusing optics, but it makes for a more reliable setup for plasma lens operation. The separable nature of the vertical and horizontal components also means that the narrow waist in the horizontal plane can be changed by replacing the final cylindrical lens f_3 without impacting the vertical plane's focusing optics. If the required plasma lens thickness is known for a given experiment, then such a cylindrical lens setup would be beneficial. While we have not yet made efforts to test such a setup at FACET-II, this would be the next step upon successful completion of the spherical focus experiment discussed in Chapter 5.

4.1.5 Ionization Refraction

One important effect that must be considered with laser ionization is the impact of how the plasma's index of refraction effects the laser pulse's propagation. Plasma has a smaller index of refraction than the neutral gas from which it is generated from. This would not be an issue if the laser pulse could be considered as instantaneous in time, but in reality a laser pulse will have some temporal width and so the laser does interact with the plasma it generates. If we think about breaking this pulse up into different slices in time, then the ionization rate in the neutral gas can be calculated for each temporal slice in the laser pulse. What can then happen is if the front of the laser is capable of ionizing a plasma, then later temporal slices of the laser will be propagating through that plasma and experience defocusing due to the change in index of refraction. This will cause those later slices of the laser to refract away from the idealized focus. The result is a wider focus region that, for cases with high refraction, can drop the intensity down to the point where the focal region no longer is fully ionizing the gas.

To simulate this effect, we first need to determine the index of refraction of the gas-plasma mixture to be able to model this behavior. For a linear medium, the index of refraction n is given by the polarization $\mathbf{P}$ through the first-order susceptibility $\chi^{(1)}$:

$$n^2 = 1 + \chi^{(1)} \tag{4.20}$$

where

$$\mathbf{P} = \epsilon_0 \chi^{(1)} \mathbf{E}. \tag{4.21}$$

We can use the polarization model given in Jackson Sec. 7.5 [29], which starts as treating an electron bound to some point and driven by an electric field. The equation of motion for such a system is

$$m_e(\ddot{x} + \gamma\dot{x} + \omega_0^2 \mathbf{x} = -e\mathbf{E}(t) \tag{4.22}$$

with an assumption that the oscillation amplitude is small compared to the optical wavelength. Next, assuming the laser's electric field is a linearly polarized plane wave $\mathbf{E} = \mathbf{E_0} e^{-i\omega t}$, the solution

to Eqn. 4.22 is that of a driven damped harmonic oscillator:

$$\mathbf{x} = -\frac{e}{m_e}\frac{1}{\omega_0^2 - \omega^2 - i\omega\gamma}\mathbf{E}. \tag{4.23}$$

For a single electron, the dipole moment is

$$\mathbf{p} = \frac{e^2}{m_e}\frac{1}{\omega_0^2 - \omega^2 - i\omega\gamma}\mathbf{E}. \tag{4.24}$$

This behavior signals that such a system is a linear material, in which case the first order susceptibility is

$$\chi^{(1)} = -\frac{n_e e^2}{\epsilon_0 m_e}\frac{1}{\omega^2} = -\frac{\omega_p^2}{\omega^2}. \tag{4.25}$$

Here we have assumed that $\omega_0 = 0$ because the electrons are not bound to any particular point, and that the dampening term γ is negligible. For a plasma, the index of refraction is then written as

$$n = \sqrt{1 - \frac{\omega_p^2}{\omega^2}}. \tag{4.26}$$

Next we calculate the index of refraction within the neutral gas. In practice, the dipole moment from every electron within an atom needs to be calculated and the sum of these dipole terms is the atomic polarizability $\alpha(\omega)$. For a gas of number density n_g, the polarization is written as

$$\mathbf{P} = \epsilon_0 n_g \alpha(\omega)\mathbf{E} \tag{4.27}$$

and the index of refraction is

$$n = \sqrt{1 + n_g\alpha(\omega)}. \tag{4.28}$$

Taken together, Eqns. 4.26 and 4.28 give the index of refraction for a gas-plasma mixture with the number densities of n_e and n_g, respectively:

$$n = \sqrt{1 + n_g\alpha(\omega) - \frac{\omega_p^2}{\omega}}. \tag{4.29}$$

Now, for relevant parameters of laser wavelength, plasma density, and gas species, Eqn. 4.29 is close to 1 and so we can use the expanded approximation for the total index of refraction:

$$n(\omega) \approx 1 + \frac{1}{2}n_g\alpha(\omega) - \frac{n_e e^2}{2\omega^2\epsilon_0}. \tag{4.30}$$

With the index of refraction known, we can now go about simulating a laser pulse propagating through a gas-plasma mixture and ionizing it as it goes. We simulate this process numerically using the split step Fourier algorithm outlined in Ref. [33]. The evolution of a laser propagating in a medium is found by solving the following set of differential equations:

$$\nabla^2 U - \frac{n^2(\mathbf{x})}{c^2}\frac{\partial^2 U}{\partial t^2} = 0 \tag{4.31}$$

$$\frac{dn_g(\mathbf{x})}{dt} = -n_g(\mathbf{x})\omega(|U(\mathbf{x},t)|). \tag{4.32}$$

Here, $n(\mathbf{x})$ is related to n_g through Eqn. 4.30 and $n_e = n_0 - n_g$. The variable U represents the scalar approximation to the wave equation,

$$\left(\nabla^2 - \frac{n^2}{c^2}\frac{\partial^2}{\partial t^2}\right) U = 0, \tag{4.33}$$

and can represent any of the three components of electric field $\mathbf{E}$. The density n also has a non-trivial spatial dependence on electric field E through the ADK model of Eqn. 4.13. We also can write the index of refraction as a constant with a perturbation, as the index of refraction is close to 1:

$$n = n_h + n_{ih} = n_h - \frac{1}{2}(n_0 + n_g)\alpha - \frac{n_e e^2}{2\omega^2 \epsilon_0}, \tag{4.34}$$

where $n_h = 1 + \frac{1}{2} n_g \alpha$.

The algorithm is as follows. First, the wave equation can be broken into two pieces from the nature of Eqn. 4.34. One piece has a constant index of refraction, and the wave equation can be solved using Fourier optics techniques such as those presented in Goodman [82] and Ratcliffe [83]. The second piece has the index of refraction as a perturbation with spatial dependence. This is simulated by propagating the laser pulse through both Fourier space and real space. This is iteratively simulated numerically using Fourier transforms:

$$U(x,y,z+\Delta z,t) = \mathscr{F}_{txy}^{-1}\left\{\mathscr{F}_{txy}\left[U(x,y,z,t)e^{ik_0 n_{ih}(t,x,y,z+\Delta z/2)\Delta z)}\right] e^{ik_z \Delta z}\right\}, \tag{4.35}$$

where k_z is a function of ω, k_x, and k_y. To include the plasma ionization in this algorithm, we simplify Eqn. 4.35 by ignoring dispersion and removing time slices t from the Fourier transforms:

$$U(x,y,z+\Delta z,t) = \mathscr{F}_{xy}^{-1}\left\{\mathscr{F}_{xy}\left[U(x,y,z,t)e^{ik_0 n_{ih}(t,x,y,z+\Delta z/2)\Delta z)}\right] e^{ik_z \Delta z}\right\}. \tag{4.36}$$

Now, $U(x, y, z, t_0$ can be solved for using $n_{ih}(t_0)$. This allows for Eqn. 4.32 to be integrated to find $n_{ih}(t_1)$, and subsequently for $U(x, y, z + \Delta z, t_1)$ to be calculated. This algorithm is illustrated through a diagram in Fig. 4.9.

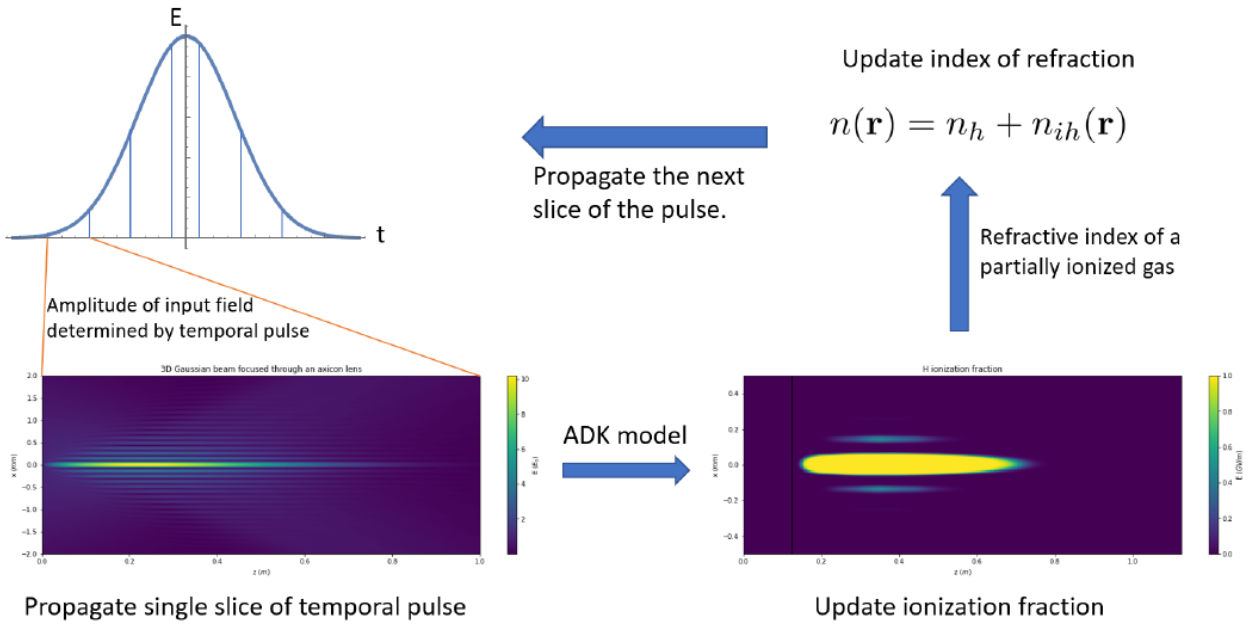

Figure 4.9: Visual diagram of the split-step Fourier propagation algorithm to numerically simulate a laser pulse ionizing a plasma. The laser pulse is divided into temporal slices, where earlier slices propagate through the gas and ionize the plasma. Later temporal slices refract off of the ionized plasma density. Reproduced from Ref. [33]

4.1.6 Crossed Cylindrical Lenses with Refraction

To demonstrate the impacts of ionization refraction on plasma lens generation, we consider the case of the crossed cylindrical lenses focus presented in Sec. 4.1.4. Ionization refraction is simulated using Fourier propagation as described in Sec. 4.1.5. The input electric field and phase profile is obtained from the complex beam parameter q of Eqn. 4.15 at a position upstream of the focus. The start of the Fourier propagation simulation is chosen to be as close to the focus as possible, but far enough away such that the intensity is too low to ionize the diatomic Hydrogen gas H_2. The density of the neutral H_2 is assumed to be uniform throughout the entire domain.

The effects of refraction can be seen in Fig. 4.10, which presents the final plasma density

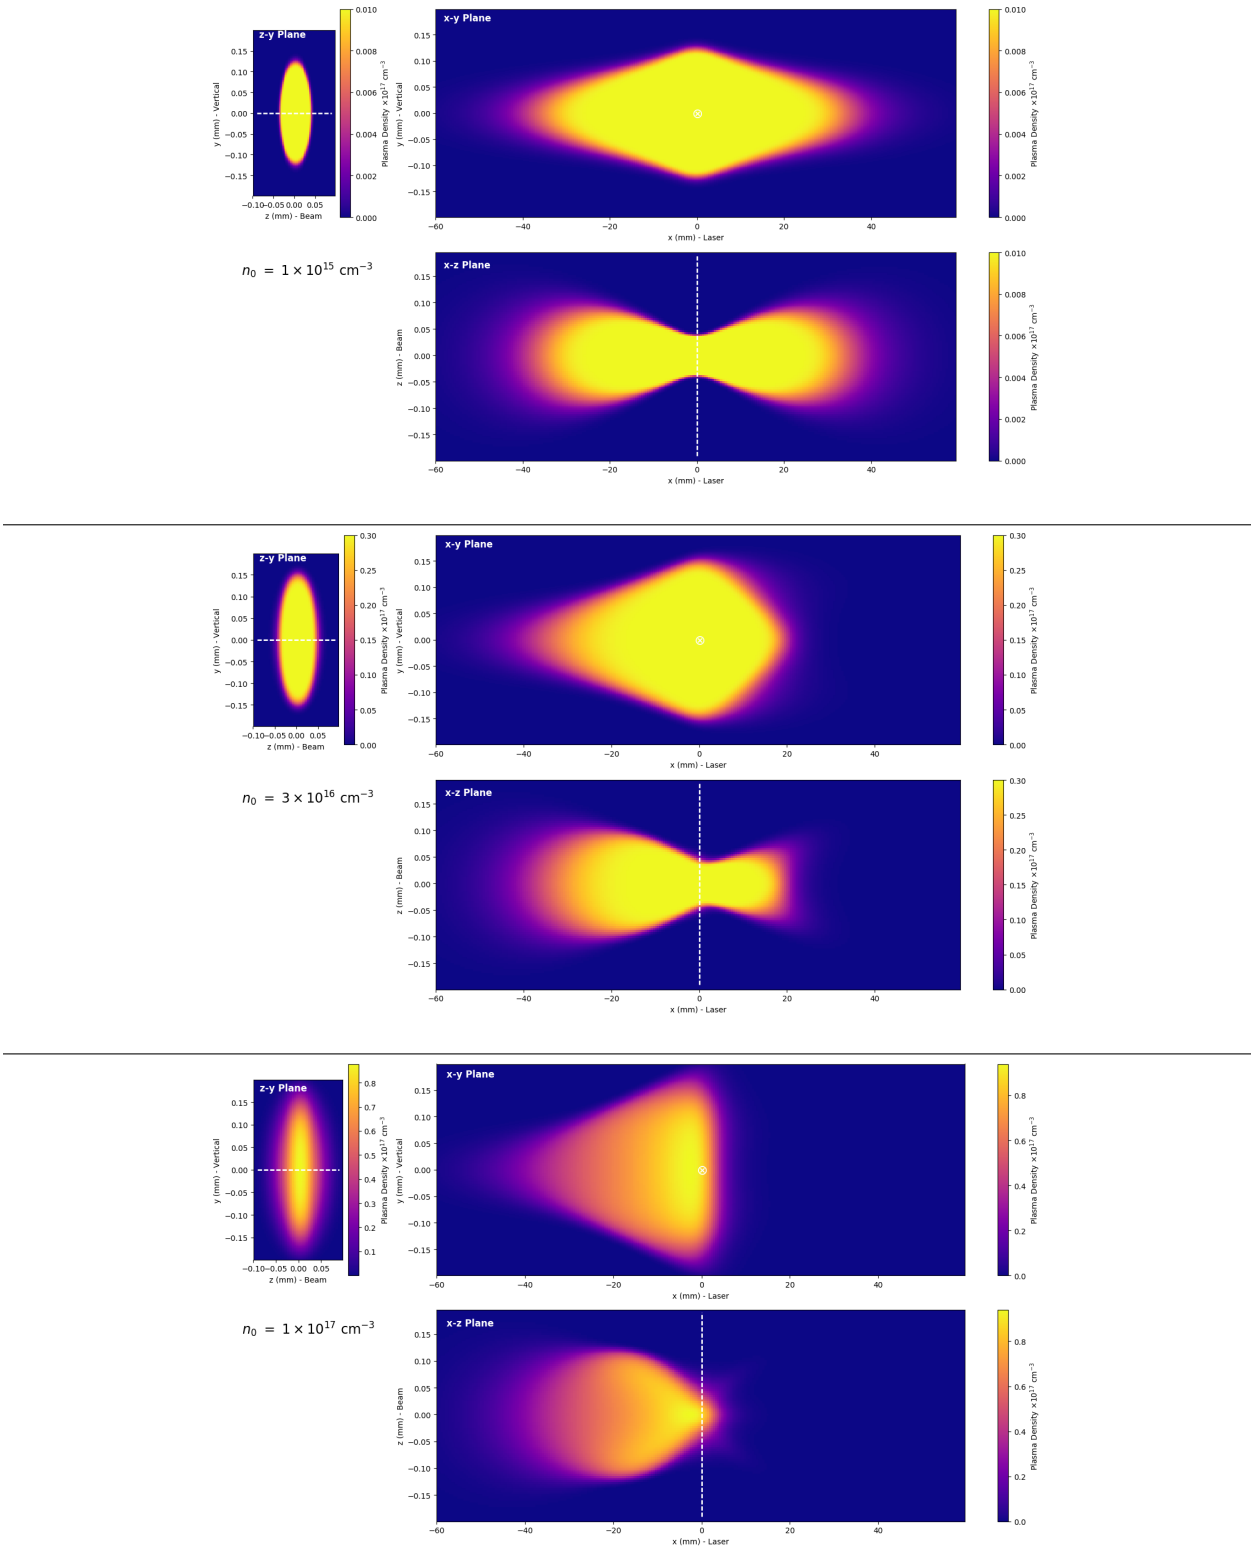

Figure 4.10: Impact of laser refraction in a uniform gas density for the crossed cylindrical lenses focus at three different densities: 10^{15} cm^{-3} (top), 3×10^{16} cm^{-3} (middle), and 10^{17} cm^{-3} (bottom).

profiles along each plane for cases with three different neutral case densities: 10^{15} cm^{-3}, 3×10^{16} cm^{-3}, and 10^{17} cm^{-3}. For a hypothetical plasma lens, the electron beam axis is shown as the dashed white lines. At the lowest density, refraction plays little to no role in altering the ionization profile. At $n_0 = 3 \times 10^{16}$ cm^{-3}, the plasma formation has refracted the laser to the point where the intensity profile downstream of the focus is significantly weaker. However, at the focus the density profile is still uniformly fully ionized and only slightly broadened. Once at density $n_p = 10^{17}$ cm^{-3}, plasma formation upstream of the focus is able to refract the ionizing laser to the point where the gas at the laser focus is no longer fully ionized.

In the context of a plasma lens, the lower two density cases yield a uniform density profile that would be suitable for ideal plasma lens focusing. Lineouts of the plasma density are shown in Fig. 4.11 for the longitudinal, electron beam axis and in Fig. 4.12 for the transverse axis. This sort of analysis gives an upper limit for the range of acceptable neutral gas densities. Then, depending on the electron beam parameters of the drive beam, a lower limit is inferred from the width of the plasma blowout. In this example, a plasma density of 3×10^{16} cm^{-3} could be used if the electron beam parameters are such that the blowout wake is no larger than 250 μm wide, as seen in Fig. 4.12.

4.2 Gas Jet Analysis

Now that we have a description for the ionization profile, we move to considering the density profile of the neutral gas we are ionizing. In Sec. 4.1.6 we made the assumption that the gas density is uniform and completely fills the experimental area. This results in a plasma density profile given by the laser and the effects of ionization refraction, and can generate suitable plasma lenses with the correct parameters. However, while this setup is sufficient for generating the plasma lens it is less than ideal for the electron beam. An electron beam propagating through non-vacuum will have its quality degraded through collisions, and the electron beamline itself is often kept under vacuum when in operation for a variety of reasons. We also saw in the previous section that the ionizing laser refracts when ionizing gas upstream of the designated plasma lens cross-section. A gas jet which provides a localized gas outflow minimizes the effect of refraction. Additionally, a

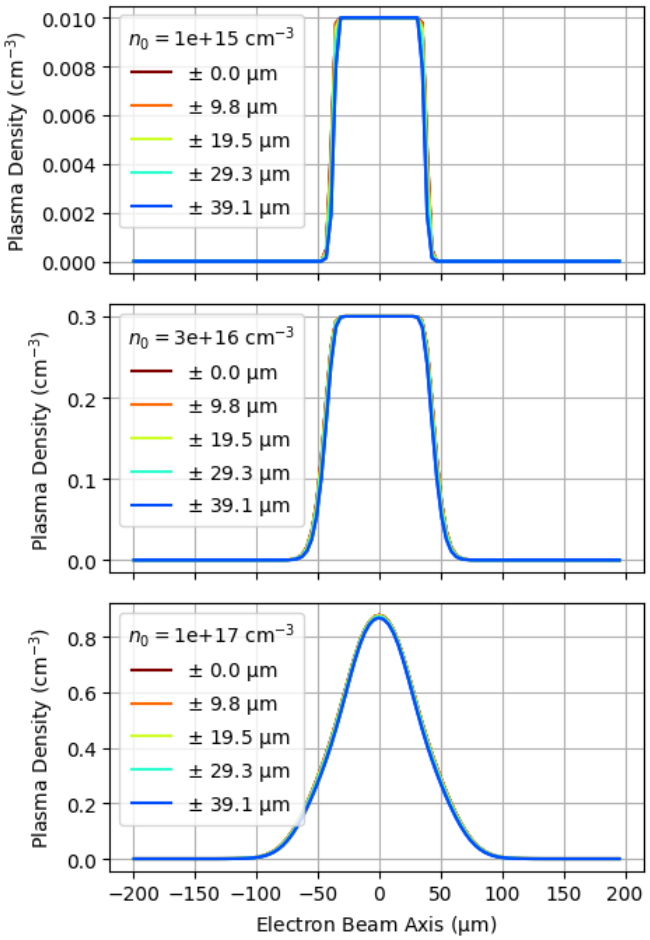

Figure 4.11: Slices of the plasma density profile in the longitudinal (electron beam) axis for three different density cases. In each subplot, the red to blue curves represent transverse offsets in the vertical y axis, and here they are nearly uniform for at least$\pm$ 40 μm. While the lower two density cases are uniformly ionized, the high density case is not fully ionized due to refraction.

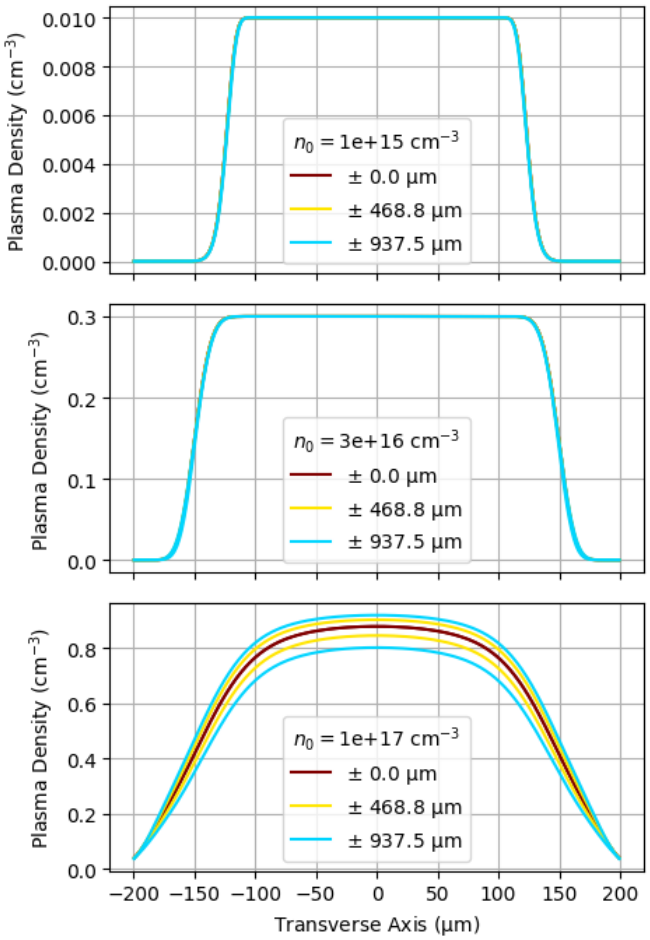

Figure 4.12: Slices of the plasma density profile in the vertical transverse axis for three different density cases. In each subplot, the red to blue curves represent transverse offsets in the horizontal x axis, for which the lower two density cases have uniformity for at least $\pm$ 940 μm. At its narrowest, the two lower density cases still have a plasma width of 200 and 250 μm, respectively. The higher density case, however, exhibits non-uniformity in both transverse axes, thus making it unsuitable for ideal plasma lens operation.

volume of neutral gas after the plasma lens would risk beam ionization as the electron beam comes to a focus. For these reasons, in this section we want to consider other methods to produce a more spatially confined gas source.

A natural choice for the gas source of a plasma lens is in the outflow of a gas jet. Gas jets themselves can come in many shapes and sizes, but if we consider a typical nozzle then we can expect gas outflows that are on the order of a few mm's wide. This is sufficient if we are laser-ionizing a plasma lens that is high density and requires a small footprint in the beamline. One could even consider more exotic designs using gas jets, such as using multiple gas jets which produce a combined gas outflow of a specific profile or using a gas cell which confines the gas outflow with solid boundaries to have better control over the density profile. Here, we only consider the simple case of a single gas jet directed upwards in y towards the axis of laser ionization x and electron beam propagation z.

The main downside with using a gas jet rather than simply ionizing a static volume of gas is that the uniformity of the gas density must be examined. In this section, we look at the density profile for a simple gas jet with typical dimensions using a fluid simulation. We then take the density profile and propagate the crossed-cylindrical laser focus through the gas jet outflow to examine the plasma lens density profile with refraction effects.

4.2.1 Conical Nozzles

Here we look at one of the most straightforward nozzle designs: that of a conical nozzle. These nozzles have circular cross sections and are axisymmetric about the upwards, axial axis. The gas jets operate in a pulsed operation mode, where the pressurized gas line that feeds the nozzle is shut off for a majority of the time. The entrance to the nozzle can be controlled by a solenoid valve that has opening times of a few ms for a single shot. Between shots, the gas diffuses around the vacuum chamber and is pumped out through vacuum pumps that return the vacuum chamber to pre-shot vacuum conditions.

Gas jet outflows and their density profiles have been studied previously by using Mach-

Zehnder interferometery and Abel inversion techniques that use the phase shift of a probe laser to calculate the density profile [84]. For typical conical nozzles operating into a vacuum environment, one can expect the outflow's density profile to have an exponential-like drop-off in the axial, outflow direction. The gas also expands radially outward from the axial axis, resulting in a Gaussian-like profile in this radial direction away from the nozzle. For minimal aberrations in beam-plasma interactions, a uniform plasma density profile is desired. However, since the size of the plasma blowout bubble is on the order of $100's$ μm wide, the density profile of the gas jet outflow is often changing on a length scale much larger than the length scale of the plasma blowout. We will revisit this later in Sec. 4.3.

Conical nozzles can also be designed with a narrow throat that diverges to the nozzle opening [85]. This type of supersonic nozzle is similar to the de Laval nozzle used in rockets, where the conversion of pressure to highly-directed flow is of key importance [86]. In these nozzles, the narrow throat results in the flow of gas increasing to the speed of sound at the throat. Once above the speed of sound, when the nozzle radius diverges, the speed of the gas further increases to several Mach numbers.

While this is useful in the design of rockets that need thrust, it is still useful when designing gas jets that operate stationary in a vacuum chamber. A more directed outflow in the axial direction means there is less radial flow of gas away from the axial axis, which in turn means that the density of the gas outflow is more uniform above the nozzle opening. However, it is impossible to have a perfectly directed gas outflow when the gas jet is opening out into a vacuum environment. There will always be some degree of radial outflow and density nonuniformity in a gas jet outflow in vacuum, so while the nozzle can be designed to minimize this effect we still need to analyze if the density profile in a typical gas jet outflow is uniform enough for the purposes of a plasma lens.

4.2.2 OpenFOAM Fluid Simulation

In addition to other studies on gas jet outflow density profiles, we perform our own simulations to have 3D density distributions that we can use in tandem with 3D laser intensity profiles for

ionization modeling. As noted in Sec. 2.4.3, the open-source, computational fluid dynamics software, OpenFOAM [60], is used to simulate a compressible Hydrogen gas propagating through a nozzle and out into a high vacuum space in 3D. For the example used here, we simulate a gas jet with a standard conical design that has a narrow throat and a long, diverging nozzle.

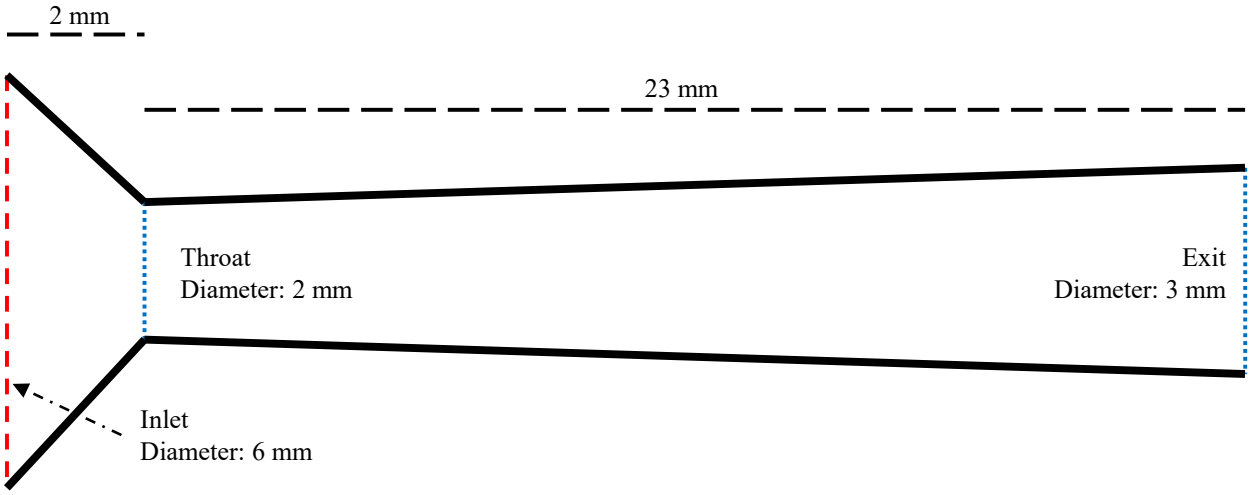

Figure 4.13: Dimensions (not drawn to scale) of the gas jet used in the OpenFOAM simulation discussed in this example. The inlet is simulated as a wall of high pressure and the rest of the domain is at near-vacuum levels. Gas passes through the throat and exit of the nozzle, and flows into a large open outflow space beyond the exit that is not pictured here.

The gas jet geometry (Fig. 4.13) is a conical nozzle with a 6 mm diameter gas inlet, 2 mm diameter throat, and 3 mm diameter nozzle exit. The distance from inlet to throat is 2 mm and the distance from throat to exit is 23 mm. The outflow region is 60 mm from nozzle exit to the top and 60 mm wide radially. The lower boundary of this outflow region is a fixed wall.

At the inlet, the initial pressure is 41430 Pa, or approximately 10^{19} cm^{-3}. The inlet boundary condition is set to 'totalPressure' which allows the pressure and gas velocity to evolve at the boundary. The background pressure is set to 0.04143 Pa, or approximately 10^{13} cm^{-3}, everywhere else. The outlet boundary conditions are 'waveTransmissive' and are set to an infinite field of 0.04143 Pa 1 m beyond the boundary. Pressure fields have specific heats set to $\gamma = 1.67$. Along the fixed walls, the pressure has a zero gradient boundary condition.

The temperature is initially set to 300 K everywhere, with a fixed value at the inlet and an 'inletOutlet' boundary condition for the outlets and fixed walls. The velocity is initially set to zero everywhere, with a 'pressureInletOutletVelocity' boundary at the inlet to match the pressure's inlet boundary condition. The outlets are set to an 'inletOutlet' boundary condition, with a 'noSlip' boundary condition along the fixed walls.

For computational efficiency, the simulation is run at a lower gas jet pressure than would be required by the example considered in this work. The results are then scaled by a linear factor of 10.12 to produce a gas density of 3×10^{16} cm^{-3} at a location 5 mm above the nozzle exit. This factor increases to 337.3 to have a target gas density of 10^{18} cm^{-3}. This is an approximation, but in general it is safe to assume that the gas jet outflow density profile linearly scales well enough with the backing pressure.

Figure 4.14 shows a cutaway final density profile along the gas jet mid-plane. The thin plasma lens will be formed in a volume located $y_0 = 5$ mm above the nozzle. About this region for up to $\pm$ 120 mm radially and $\pm$ 0.2 mm vertically, we make an approximate fit to the radial and axial density profiles to be used in the split-step Fourier propagation algorithm of Sec. 4.1.5. In the vertical direction, the gas jet outflow follows an exponential profile of approximately

$$n_{jet}(x = 0, y, z = 0) = c_1 \exp\left(-\frac{y + y_0}{c_2}\right) + c_3; \tag{4.37}$$

$$c_1 = 3.32 \times 10^{16} \text{ cm}^{-3}; \quad c_2 = 2.00 \text{ mm}; \quad c_3 = 2.39 \times 10^{14} \text{ cm}^{-3}. \tag{4.38}$$

For the radial direction of the gas jet outflow in the x-z plane, the profile is approximately Gaussian with an rms size of $\sigma_0 = 2.99$ mm at y_0. But, as the gas outflow expands as it travels away from the nozzle exit, this rms size gradually increase with larger y. To account for this, we make a linear fit for the vertical evolution of the radial Gaussian rms size:

$$\sigma(y) = \sigma_0 + \sigma_{(1)} y + \sigma_{(2)} y^2; \tag{4.39}$$

$$\sigma_0 = 2.99 \times 10^3 \ \mu\text{m}; \quad \sigma_{(1)} = 0.550; \quad \sigma_{(2)} = -5.09 \times 10^{-6} \ \mu\text{m}^{-1}. \tag{4.40}$$

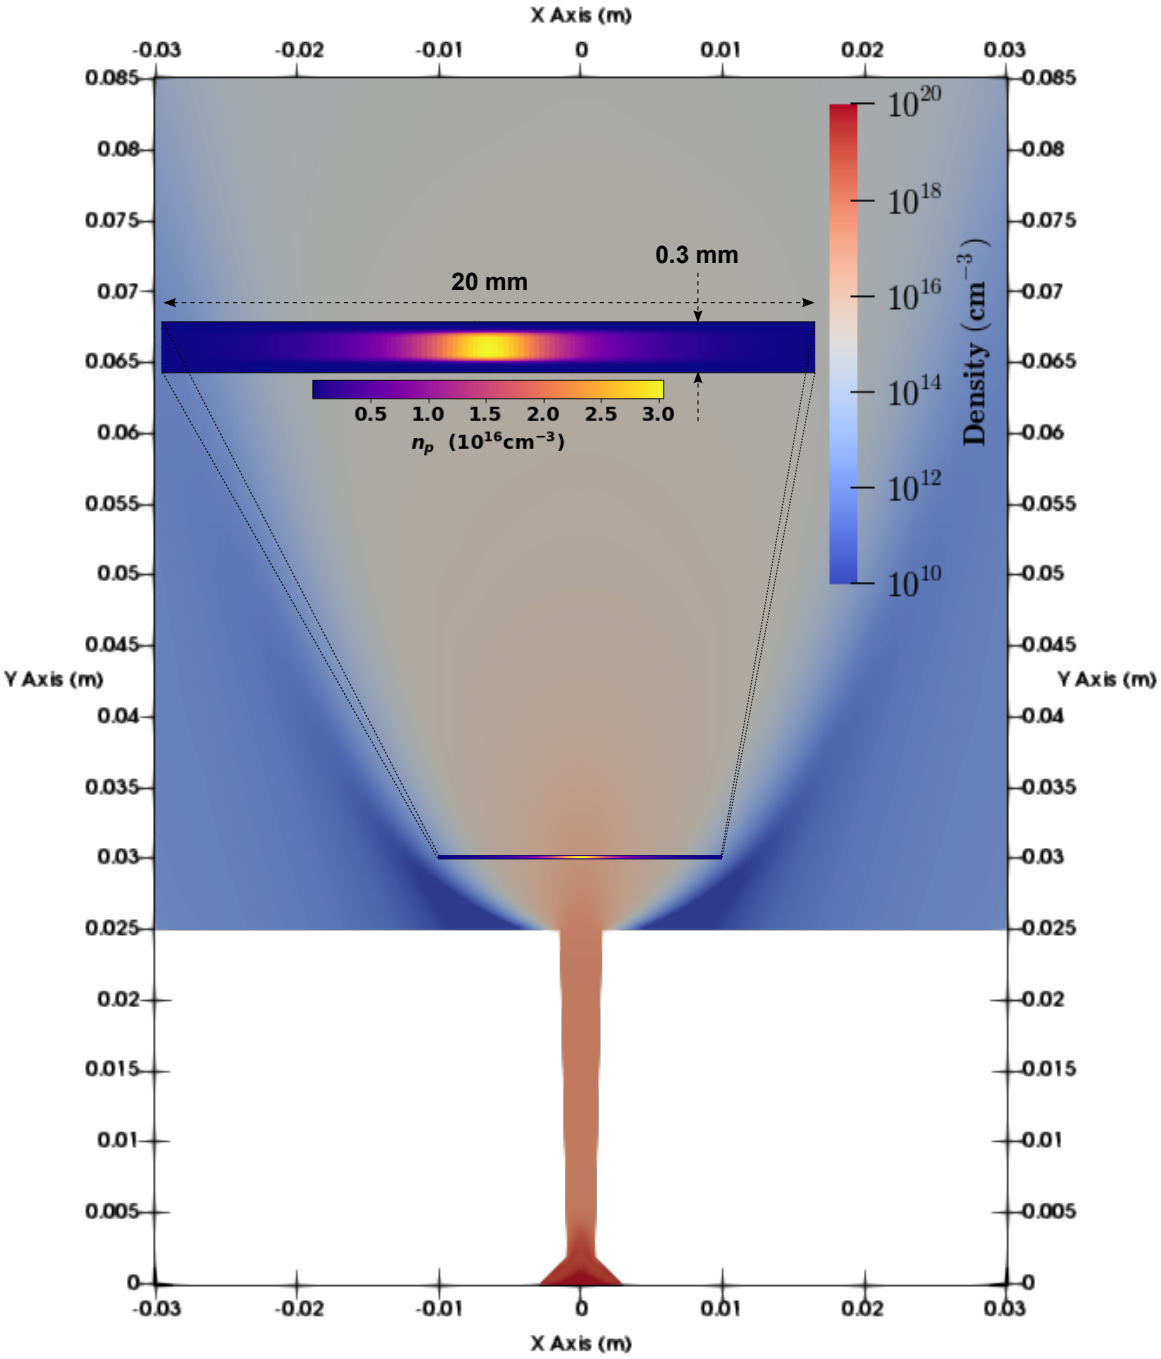

Figure 4.14: Gas jet density profile at the mid-plane of a 3D OpenFOAM simulation at time step 140 μs. Insert: magnification of laser-ionized plasma density profile. An electron beam's trajectory would be into the page.

Using this vertically-varying rms size, the full empirical density profile for the gas jet outflow is

$$n_{jet}(x,y,z) = \exp\left(-\frac{x^2+z^2}{2\sigma(y)}\right)\left[c_1\exp\left(-\frac{y+y_0}{c_2}\right)+c_3\right]. \tag{4.41}$$

In the context of density uniformity, if we consider that the length scales of a plasma blowout are on the order of 100 μm then it is clear that the Gaussian rms size of the x-z profile is much larger. Thus, we can safely assume that the the gas density in these axes is uniform and the plasma density profile is entirely given through the laser ionization. For the vertical y axis, taking the derivative of Eqn. 4.41 at the origin (in other words, centered at y_0) reveals a vertical density gradient of 1.38×10^{13} cm$^{-3}/\mu$m when the target density is scaled to 3×10^{16} cm^{-3}. For a blowout of width 100 μm, this is a density variation of 4.6% across the transverse plasma wake width. The exponential profile has a length scale of $c_2 = 2.00$ mm so this is gradient is approximately constant over the width of the plasma blowout.

4.2.3 Crossed Cylindrical Lenses with Gas Jet Density Profile and Refraction

Now that we have a gas density profile, we can use this in conjunction with the split-step Fourier propagation algorithm of Sec. 4.1.5 to predict the final plasma density profile. Our procedure here is the exact same as in Sec. 4.1.6 that uses the laser profile of the crossed cylindrical lenses from Sec. 4.1.4, except now we need to account for a nonuniform initial gas density given by Eqn. 4.41.

Since the neutral gas is more spatially contained when using a gas jet, the effect of refraction will be much more limited. While the laser intensity is high enough to ionize gas within a region of $\pm$ 40 μm (Fig. 4.8), the gas jet outflow only exists in a smaller region. Therefore, we can use a high number density in the gas jet outflow than in the uniform gas density examples of Sec. 4.1.6. Here, we propagate the laser through three different gas jet outflow densities by linearly scaling the full density profile such that the number density at the origin is 3×10^{16} cm^{-3}, 10^{17} cm^{-3}, and 10^{18} cm^{-3}. The ionization profiles of these simulations are summarized in Fig. 4.15.

We can observe from Fig. 4.15 that the effects of refraction are indeed much smaller than previously in Sec. 4.1.6 and Fig. 4.10. While the examples without a gas jet outflow profile showed refraction to significantly lower the ionization rate in the center of the plasma lens at density 10^{17} cm^{-3}, using a gas jet allows for the plasma lens density to be larger than 10^{17} cm^{-3} and even

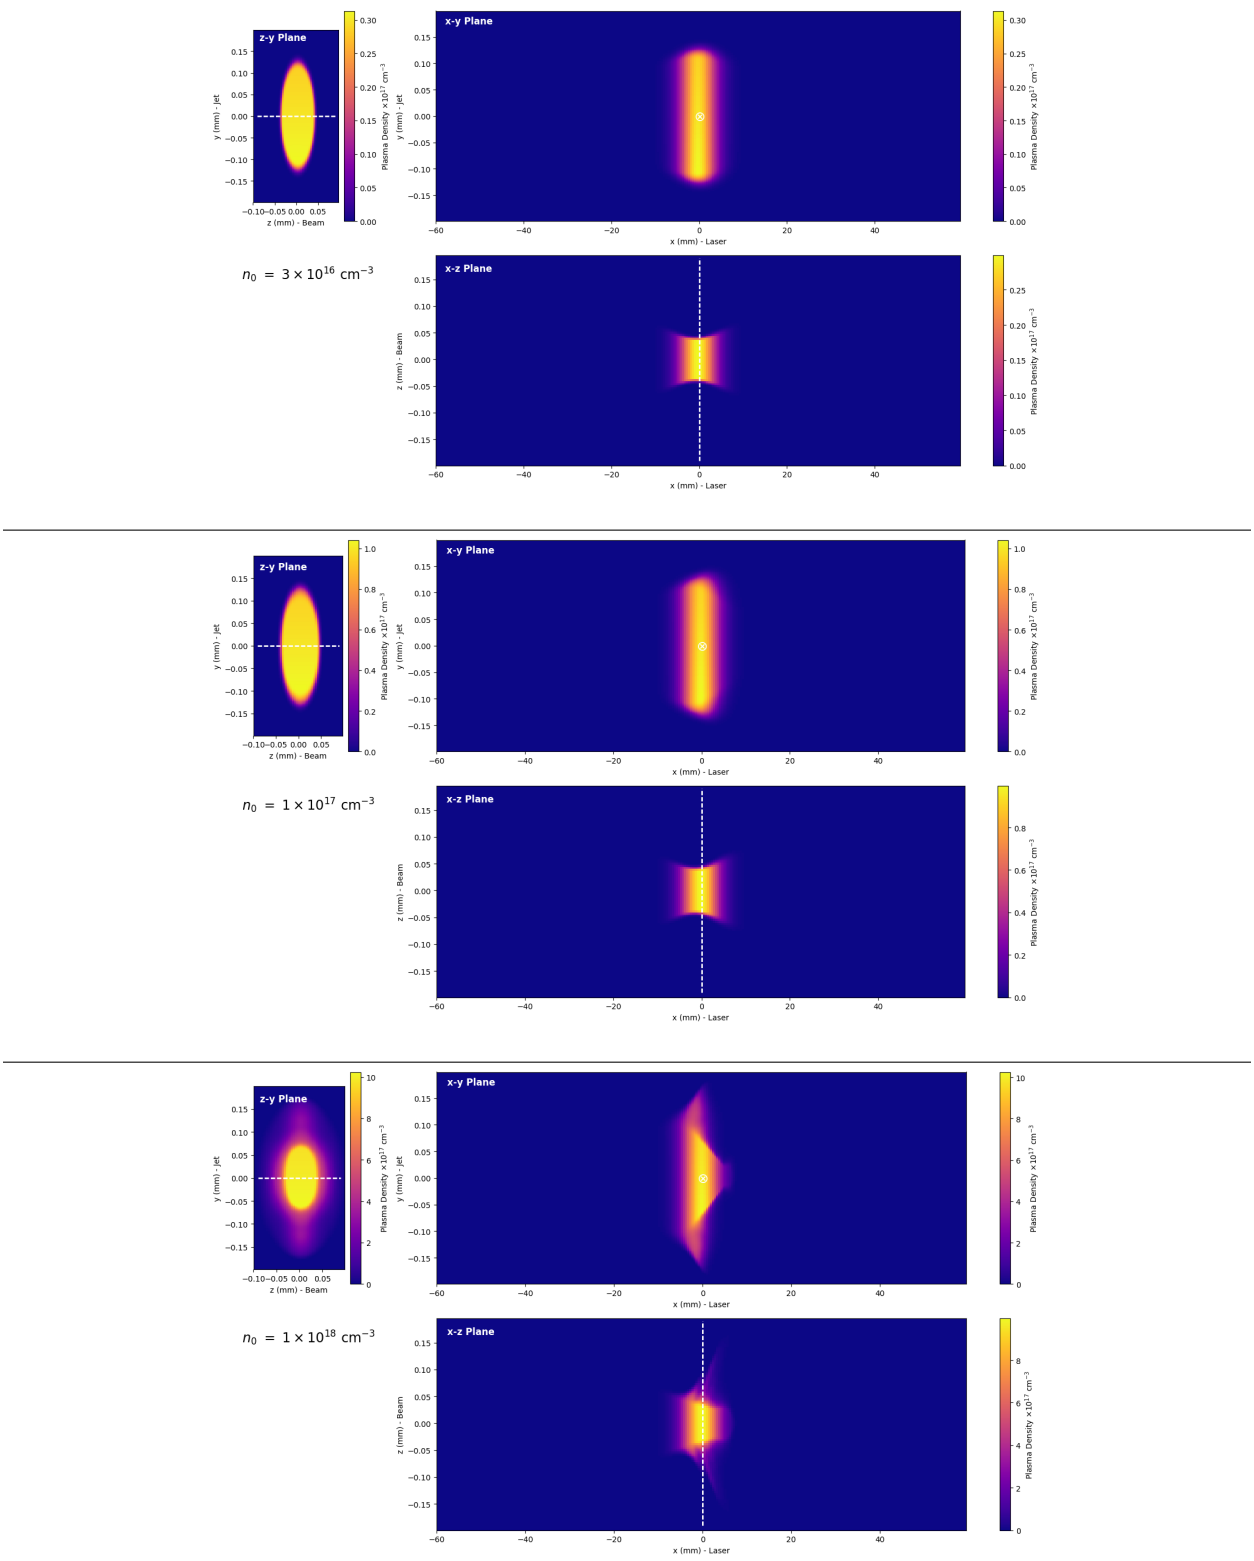

Figure 4.15: Impact of laser refraction in a gas jet outflow for the crossed cylindrical lenses focus at three different densities: 3×10^{16} cm^{-3} (top), 10^{17} cm^{-3} (middle), and 10^{18} cm^{-3} (bottom).

almost be able to operate at 10^{18} cm^{-3}. Slices of the plasma density profile are plotted for the longitudinal axis with offsets in the vertical direction (Fig. 4.16), and for the transverse vertical axis with offsets in the horizontal direction (Fig. 4.17).

One detail that can be easily noted from Fig. 4.17 is that while the central plasma density profile is fully ionized, it is not uniform. This is from the gas jet outflow density, where in the vertical direction the density follows an exponential drop-off. This arises as an approximately linear density gradient in which the plasma density changes by a few percent across the width of the nonlinear blowout. While this is a perturbation on the ideal plasma lens operation at uniform density, the consequences ultimately depend on how large of a gradient the density profile has. The impacts of such a density gradient is explored in further detail in Sec. 4.3.

Overall, if we are designing a plasma lens for an experimental setup then a gas jet is preferable so long as the density gradient is small enough. Less plasma refraction means that the plasma lens can be generated in a higher gas density for stronger focusing strengths. Depending on the experiment, it may also be easier to place a gas jet in the electron beam linac than it is to fill a chamber with uniform, dense gas. In either case, the maximum gas density is determined by the ionizing laser optics, which are in turn determined by both the plasma blowout width and required plasma lens width.

4.3 Nonuniform Plasma Densities

Up until this point, we have been assuming the usual nonlinear plasma blowout wake formation outlined in Sec. 2.3.3. A core assumption in that model is that the plasma density is uniform.[1] However, as we saw earlier in Sec. 4.2.2, a gas jet outflow inherently has a nonuniform density. If we are attempting to design a passive plasma lens using a gas jet, then we must learn how the inclusion of such a realistic density profile effects the transverse wakefields of the blowout wake. This section presents the work we have published in Ref. [87] on the topic of a linear density gradient in

[1] Or, more accurately, that the plasma density is cylindrically symmetric. If the density is a function of r, then a modified formalism can be used. See Refs. [52, 53].

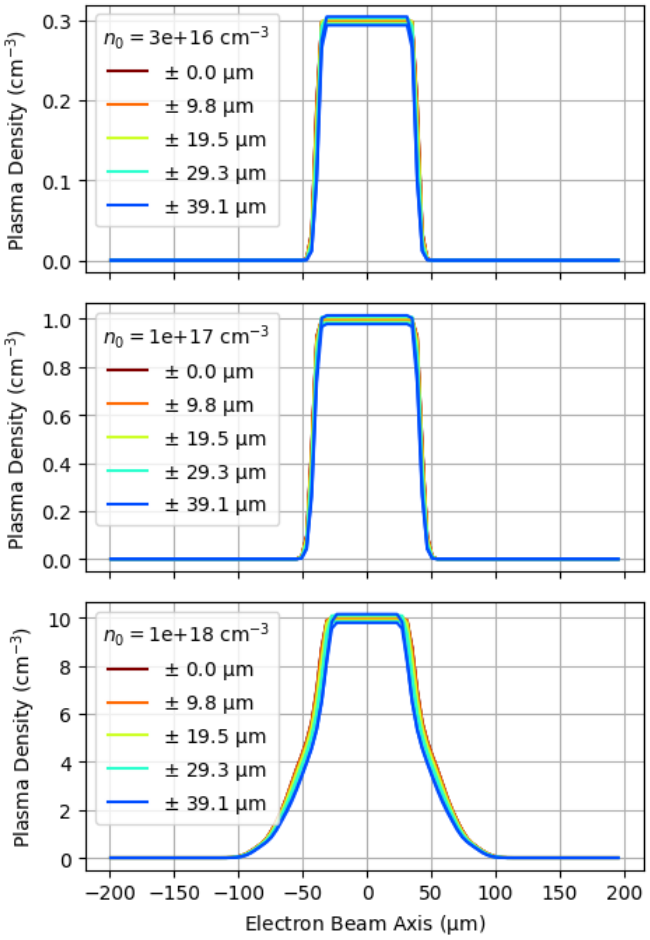

Figure 4.16: Slices of the plasma density profile in the longitudinal (electron beam) axis for three different gas jet outflow density cases. In each subplot, the red to blue curves represent transverse offsets in the vertical y axis, and exhibit a slight gradient due to the gas jet outflow density profile. While the lower two density cases are fully ionized, the high density case exhibits some refraction and has a less sharp ionization profile.

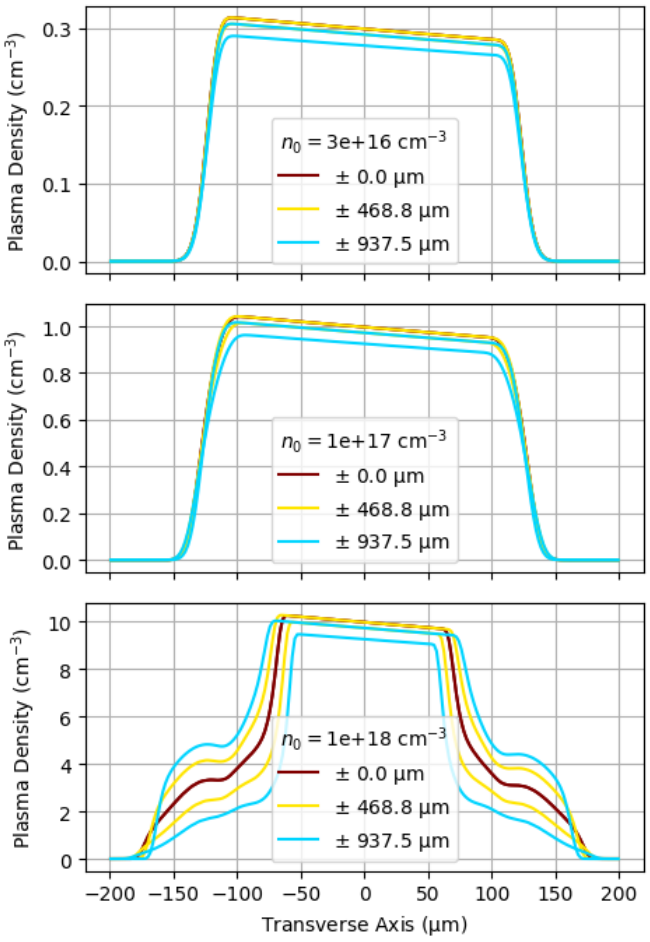

Figure 4.17: Slices of the plasma density profile in the vertical transverse axis for three different gas jet outflow density cases. In each subplot, the red to blue curves represent transverse offsets in the horizontal x axis, where there is a slight change due to the radial Gaussian density profile of the gas jet outflow. The two lower density cases have a fully-ionized plasma width of 200 μm, with a linear gradient due to the gas jet outflow. The higher density case exhibits laser refraction, which causes the fully-ionized region to be smaller at 100 μm.

a plasma lens.

We consider an underdense passive plasma lens implemented by laser ionization of a gas jet outflow, such as in Refs. [26, 27, 88]. A femtosecond laser pulse is capable of ionizing a region of plasma localized within the outflow of a gas jet [75]. If the laser focus is much smaller than the gas outflow volume, the shape of the plasma lens will be characterized by the laser parameters and focusing optics. Within this region of ionized gas, for a typical gas jet outflow the density profile would exhibit some degree of non-uniformity. A sketch of the hypothetical experimental layout we are considering is presented in Fig. 4.18. Transverse density gradients have been explored previously in the linear plasma wake regime as a pathway to a plasma-based undulator [89], but in this study we concentrate on the nonlinear blowout regime.

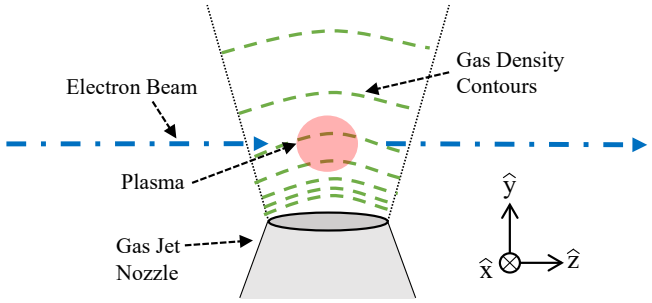

Figure 4.18: Sketch of gas jet, plasma (red region), and electron beam (dashed blue) layout. Gas jet outflow (dashed green) is upwards in the y direction with an exponential density function. The outflow in the radial, x-z plane has a smoother Gaussian profile. The pre-formed plasma is ionized within a smaller region of this gas jet outflow by a laser propagating in the x direction (not shown). The electron beam propagates left to right in the z direction.

An electron beam in the blowout wake of such a plasma lens will have a longitudinally-varying deflection along the gas jet outflow axis that scales with the magnitude of the plasma density gradient. Though there may be useful ways to exploit this behavior, in general it is likely to produce uncorrectable aberrations in the focusing of the electron beam. It is therefore important to be able to model and understand this behavior in order to sufficiently mitigate it when designing a plasma lens.

A typical density profile for a gas jet outflow can be described by a steep exponential or

Lorentzian axial descent and a more gradual Gaussian or flattop profile in the outward radial dimension [84, 85]. Here, using typical spatial coordinates for a electron beam propagating in the z direction, we assume the gas flow is directed vertically in the y direction and diffuses radially outward in the x-z plane. If we consider a gas number density 5×10^{16} cm^{-3} just above the nozzle, the corresponding plasma wavelength is 149 μm, which sets the transverse scale of the plasma wake. In comparison, Ref. [84] measured the characteristic length scale of gas jet density variation to be larger than 550 μm in the vertical, outflow axis. The radial, Gaussian length scale is also larger than the plasma wavelength and, for a plasma centered within the gas jet outflow, the density gradient in this x-z plane is negligible. While this gas jet profile was measured in a high density gas jet, gas jet outflows at lower densities can be measured using the technique described in Ref. [90]. Thus, the vertical, axial density profile can be approximated as varying linearly across the wake and the radial density profile can be assumed constant across the wake.

We quantify the transverse density gradient with the dimensionless parameter

$$g \equiv \left(\frac{\partial n}{\partial y}\right) \frac{c}{\omega_{pe} n_0}. \tag{4.42}$$

where c/ω_{pe} is the plasma skin depth, n_0 is unperturbed plasma number density at the height of the blowout center (defined as the drive beam axis) $y = 0$, and $\partial n/\partial y$ is the density gradient at $y = 0$ and is considered to be explicitly constant for the remainder of this work. Equation 4.42 quantifies the relative change in plasma density across the plasma skin depth.

For sufficiently small gradients where $g \ll 1$, we assume that the effects of the nonuniform plasma density are perturbative and the shape of the blowout wake retains its circular transverse cross-section. We will revisit the region of validity for this assumption later in Sec. 4.3.1. The density profile of the ion column is then given by

$$n_p(y) = n_0 + \left(\frac{\partial n}{\partial y}\right) y \tag{4.43}$$

over a circular cross section of blowout radius R, $x^2 + y^2 < R^2$. In this paper, we assume $\partial n/\partial y < 0$ to reflect a typical experimental setup where the gas jet is located below the electron beamline, yielding higher density at lower values of y.

Table 4.1: Single-bunch simulation parameters and results. The top sub-table lists the relevant wake dimensions for the zero-gradient cases and the input parameters across all simulations of the same density. The bottom sub-table lists the measured values of wake dimensions and perturbative magnitudes from simulations with a non-zero density gradient. Measured values are defined later in the text. All simulations have appropriate simulation size and grid scales, and drive beam parameters are chosen to produce blowout wakes in the different densities. Simulation parameters constant across all runs are listed in the text.

n_0 (cm^{-3})	R_{max} (μm)	ξ_{max} (μm)	Q_{drive} (nC)	$\sigma_{z,drive}$ (μm)	$\sigma_{r,drive}$ (μm)	dx, dy, dz (μm)	L_x, L_y (μm)	L_z (μm)
1×10^{16}	91.91	160.4	2	32.0	7.2	1.23	400	498
2×10^{16}	76.15	112.2	2	20.8	7.2	1.20	400	352
1×10^{17}	48.83	68.0	2.5	15.6	5.9	0.5	190	191

n_0 (cm^{-3})	$\partial n/\partial y$(cm^{-4})	g	R_+ (μm)	R_- (μm)	α	$\langle \partial n/\partial y \rangle_{sh}$ (cm^{-4})	$\langle E_{sh}/R_p^2 \rangle$ (V/m^3)
1×10^{16}	4×10^{17}	0.212	98.27	86.95	0.0346	$1.17 \times 10^{18} \pm 1.43 \times 10^{17}$	$6.25 \times 10^{16} \pm 1.03 \times 10^{16}$
2×10^{16}	2.5×10^{15}	0.0005	76.171	76.137	0.000104	$1.63 \times 10^{15} \pm 1.88 \times 10^{15}$	$3.83 \times 10^{14} \pm 1.48 \times 10^{14}$
2×10^{16}	2.5×10^{16}	0.005	76.24	76.07	0.00103	$7.42 \times 10^{16} \pm 2.05 \times 10^{16}$	$4.43 \times 10^{15} \pm 4.60 \times 10^{14}$
2×10^{16}	2×10^{17}	0.038	77.09	75.02	0.00793	$5.85 \times 10^{17} \pm 6.79 \times 10^{16}$	$3.62 \times 10^{16} \pm 3.24 \times 10^{15}$
2×10^{16}	8×10^{17}	0.150	80.42	72.58	0.0309	$2.40 \times 10^{18} \pm 2.63 \times 10^{17}$	$1.38 \times 10^{17} \pm 1.26 \times 10^{16}$
1×10^{17}	2×10^{18}	0.034	49.57	47.94	0.00994	$7.93 \times 10^{18} \pm 1.37 \times 10^{18}$	$3.76 \times 10^{17} \pm 3.87 \times 10^{16}$
1×10^{17}	4×10^{18}	0.067	50.49	47.21	0.0198	$1.60 \times 10^{19} \pm 2.69 \times 10^{18}$	$7.45 \times 10^{17} \pm 7.42 \times 10^{16}$

The organization of this section is as follows: In Sec. 4.3.1 we perform particle-in-cell (PIC) simulations with and without a linear plasma density gradient to observe the difference in collective effects of the plasma wake. Section 4.3.2 derives a semi-empirical model for the perturbations resulting from the density gradient: additional terms in the ion column's wakefield, the transverse drifting of the blowout wake's centroid, and the density gradient in the electron sheath. Section 4.3.3 compares two-bunch PIC simulations of a plasma lens and linear density gradient with the observed vertical deflection and longitudinal variation from models in Sec. 4.3.2. Section 4.3.4 calculates the transverse variation of the longitudinal wakefields due to the longitudinal variation of the transverse wakefields.

4.3.1 Numerical PIC Simulations

We begin our analysis by running several 3D PIC simulations of the wake formation in a plasma with a density gradient using a single, relativistic electron driver bunch. These are compared to identical simulations with no density gradient to find a semi-empirical model that describes the effect of the linear transverse density gradient. The PIC simulation software we use is VSim [59], and the different simulation parameters used are summarized in Table 4.1. Common to all simulations, the drive beam emittance is $\epsilon_N = 3.2$ mm $-$ mrad, the drive beam energy is $E = 10$ GeV, and the energy spread of the drive beam is Gaussian with $\sigma_\delta = 1$ %. The drive beam bunch length and transverse size were sufficient to drive a nonlinear blowout wake in all cases (see Table 4.1 for exact values used). In these simulations, a moving window follows the drive beam as it propagates 5 mm in the plasma until the blowout wake stops evolving and reaches a steady state.

A quantitative analysis of the simulations is given later in Sec. 4.3.2. In this section we introduce two qualitative effects apparent from analyzing the plasma electron density and blowout sheath along the vertical (plasma density gradient) direction, shown in Fig. 4.19. The longitudinal cooardinate along the drive beam axis ξ is given by $\xi = ct - z$. First, we note that the simulations show the wake sheath trajectories in regions of plasma with higher (lower) density than n_0 behave

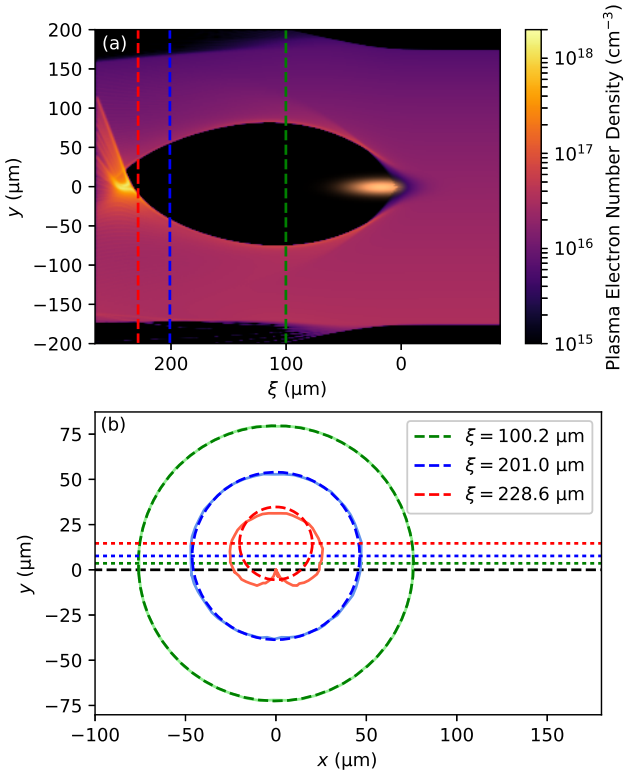

Figure 4.19: PIC simulation of a nonlinear blowout wake driven by a single electron bunch with plasma density $n_0 = 2 \times 10^{16}$ cm^{-3} and a linear, vertical density gradient of $g = 0.150$. Subplot (a) shows a slice in the longitudinal-vertical plane, with the drive beam propagating towards the right. The colorbar represents the plasma electron density, and the drive beam is plotted with an arbitrary colorbar. The ions are assumed to form a stationary background with a number density given by Eqn. 4.43. The three dashed lines mark the longitudinal positions that are examined in subplot (b), which plots the blowout wake's transverse cross-section. The solid lines are the wake boundary from simulation and the dashed lines are circular fits to the boundary. The dotted lines mark the wake's vertical center in each of the three circular fits. The black dashed line marks $y = 0$. The wake's centroid shifts upwards toward lower densities near the rear of the blowout while the cross-section becomes less circular.

as if the plasma skin depth is smaller (larger), as expected. The sheath in the lower density region will be longer and have a wider blowout radius than the sheath in the higher density region. While the cross-sectional shape of the rear of the wake is significantly altered by this dynamic, the bulk of the wake remains relatively circular. We can approximate longitudinal slices along the bulk of the blowout as having a circular transverse cross section with the center offset toward the lower density region. The magnitude of this offset grows along the length of the wake and, as a result,

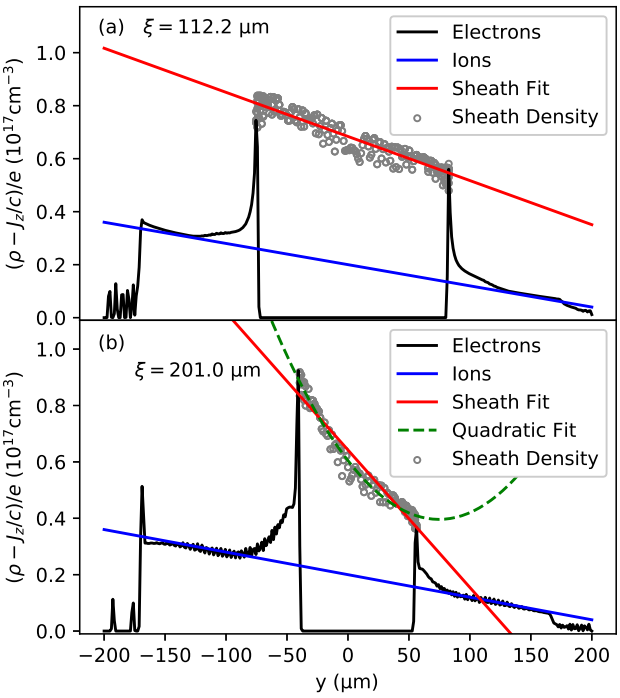

Figure 4.20: Sum of plasma electron density and axial current (solid black) along y at $x = 0$. Subplot (a) and (b) are at longitudinal positions $\xi = 112.2$ μm and $\xi = 201.0$ μm, respectively. The density and gradient in this simulation are $n_0 = 2 \times 10^{16}$ cm^{-3} and $g = 0.150$. The initial plasma density before the arrival of the beam is shown in blue, which is also the density of the stationary ions. The projection of the electron sheath's peak for all x onto the y axis is shown by the gray circles, which represent the wake boundary for a given longitudinal slice. Linear fits to the peak sheath density across y are shown in solid red. In subplot (b), towards the back of the wake the sheath density profile is not modeled perfectly by a linear function, and a sample quadratic fit is shown in dashed green

modifies the wakefield along the electron drive beam axis.

Secondly, a closer look at the blowout electron sheath reveals a linear, vertical density and axial current gradient with a separate magnitude than that of the unperturbed neutral plasma (Fig. 4.20). Here, as in Ref. [46], we consider the combined contribution of the electron charge density and axial (with respect to drive beam propagation axis) current as the full electron sheath profile:

$$n_{sh} = -(\rho - J_z/c)/e \tag{4.44}$$

The result of this non-axisymmetric electron sheath distribution contributes to the vertical com-

ponent of the transverse wakefield within the blowout wake. However, toward the rear of the wake the vertical sheath dependence becomes much steeper than in the center of the wake and, in cases with a higher plasma density gradient, the sheath profile can develop a quadratic or exponential form.

4.3.2 Analytic and Empirical Modeling

The most significant contribution to the transverse wakefield within the nonlinear blowout wake comes from the stationary, positively charged ions. While the transverse wakefield due to a uniform density ion column is linear and axisymmetric, the presence of the linear ion density gradient breaks this symmetry. To reach an analytic solution for the wakefield from an ion column described by Eqn. 4.43, we first assume that the blowout itself is circular with radius R_p. The simulations presented in Sec. 4.3.1 show that this is a valid assumption everywhere except the rear-most region of the wake.

Under the assumption that the ion column is circular for a given longitudinal slice, we find a solution to the 2D Poisson-like equation [46] which gives the following form for the potential:

$$\psi = -\pi e n_0 \left(x^2 + (y - 2\bar{Y}(\xi))^2\right) - \frac{1}{2}\pi e \frac{\partial n}{\partial y}\left(y^3 + x^2 y\right). \tag{4.45}$$

Here, e is the fundamental electric charge and

$$\bar{Y}(\xi) = \frac{1}{4}\frac{R(\xi)^2}{n_0}\frac{\partial n}{\partial y} \tag{4.46}$$

is the "center of charge" for an ion column with a linear gradient in charge density. Because the blowout radius $R(\xi)$ has a longitudinal dependence, the transverse wakefields will correspondingly have a longitudinal variation. We look at the effects of this longitudinal variation later in this section.

Using Eqn. 4.45, the transverse wakefields from the ions are calculated as $W_{x,y} = -\nabla_{x,y}\psi$:

$$W_x = 2\pi e n_0 x + \pi e \frac{\partial n}{\partial y}\left(xy\right) \tag{4.47}$$

and

$$W_y = 2\pi e n_0 (y - 2\bar{Y}(\xi)) + \frac{1}{2}\pi e \frac{\partial n}{\partial y}\left(3y^2 + x^2\right). \tag{4.48}$$

There are two notable features in Eqns. 4.45, 4.47, & 4.48: first, the $-2\bar{Y}$ term in Eqn. 4.48 indicates a transverse offset of the center of the linear focusing term towards the higher density side of the ion column. The second feature is the sextupole-like rightmost term in Eqn. 4.45. The $x-y$ plane has the same y^3 dependence that produces a y^2-dependent focusing force of a sextupole magnet [28]. Interestingly, when compared to a sextupole magnet, the plasma lens is significantly less sensitive to beam centroid misalignment in the perpendicular x-axis, but just as sensitive in the y-axis. This feature, along with the compact size and high degree of tunability of the underdense plasma lens may make it an attractive alternative to sextupole magnets for correcting higher order features in electron beams. However, it could prove difficult to access these fields since the strength of the sextupole fields are tied to the strength of the ion column's coincident focusing fields, as well as the vertical offset of the blowout center.

The effect of the fields expressed in Eqns. 4.47 & 4.48 contribute most significantly to the transverse field inside the blowout wake, but there remain two outstanding features that lead to perturbations of the transverse wakefield: the vertical shifting of the blowout wake's centroid and the asymmetric electron sheath density and current profiles. These effects are far more difficult to describe analytically, so we have adopted simulation-based empirical methods to predict their effects.

First, we model the vertical drifting of the geometric wake center. This drift is due to the relatively weaker (stronger) restorative force on the plasma electrons from the ion column on the lower (higher) density side of the plasma. The imbalance causes the sheath to close more slowly (quickly) on the lower (higher) density side, shifting the transverse geometric center of each longitudinal slice of the blowout wake more to the lower density side from the front to the rear. In subplot (a) of Fig. 4.21, the sheath profiles of the lower and higher density sides of the wake are plotted and labelled as "greater" and "lesser", respectively. For the bulk of the wake, the geometric

center drift is close to a linear function with respect to longitudinal position within the blowout wake (Fig. 4.21).

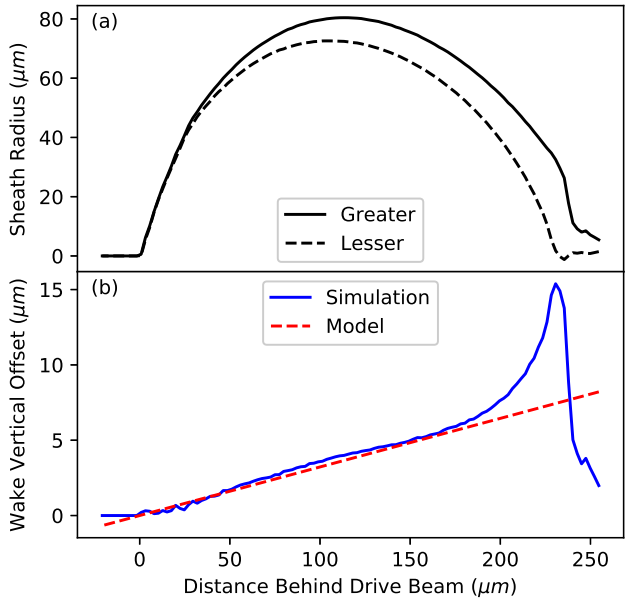

Figure 4.21: Longitudinal evolution of the blowout wake's geometry and its sheath. In subplot (a) the greater and lesser curves refer to the sheath trajectories in the lower and higher density regions, respectively. In (b) the wake's centroid from a circular fit is plotted against the empirical model from Eqn. 4.50. The model accurately predicts the wake's vertical offset everywhere except in the rear of the wake.

We start by defining R_+ and R_- as the maximum radius of the greater and lesser sheaths, respectively. An empirical model that agrees well with the simulations in Sec. 4.3.1 is that R_+ and R_- are given by

$$R_\pm \approx R_{\max}\sqrt{1 \pm \left|\frac{\partial n}{\partial y}\right| \frac{R_{\max}}{n_0} A} \quad ; \quad A \equiv 0.311 \tag{4.49}$$

where $R_{\max}$ is the maximum blowout radius in a uniform plasma density n_0 and A is an empirical constant. The square root dependence is motivated by the scaling of the wake size with plasma wavelength [46]. Assuming the offset of the geometric wake center is 0 at the front, grows linearly, and is $(R_+ - R_-)/2$ at the location of $R_{\max}$; the wake center is given by

$$y_c(\xi) \approx -\text{sgn}\left(\frac{\partial n}{\partial y}\right) \frac{(R_+ - R_-)}{2\xi_{\max}} \xi \equiv -\alpha\xi \tag{4.50}$$

where $\xi_{\max}$ is the longitudinal position of $R_{\max}$ and the scalar α is the slope of this vertical offset

shift. Equation 4.50 is designed to have the opposite sign of $\partial n/\partial y$. As shown in Fig. 4.22, the empirical constant A is found by fitting this model to the measured center offset growth across all of the simulations (Table 4.1) in the longitudinal region

$$\frac{1}{2}\xi_{\max} < \xi < \frac{3}{2}\xi_{\max} \tag{4.51}$$

where the growth is linear, as demonstrated in Fig. 4.21.

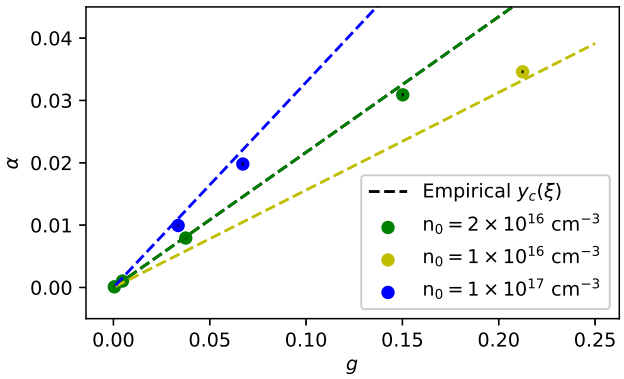

Figure 4.22: Comparison of vertical offset growth along ξ observed in simulations (dots) with the theoretical prediction of Eqns. 4.49 and 4.50 (dashed lines). The simulations are grouped according to central density n_0, as $R_{\max}$ and $\xi_{\max}$ both depend on n_0.

This offset of the wake center modifies our previous equations slightly, assuming that we want to keep the origin anchored to the drive beam axis. The y coordinate must now be expressed as $y - y_c(\xi)$, giving the following expressions for the transverse wakefields inside the wake:

$$W_x = 2\pi e n_c(\xi) x + \pi e \frac{\partial n}{\partial y}\Big(x(y - y_c(\xi))\Big) \tag{4.52}$$

and

$$W_y = 2\pi e n_c(\xi)(y - y_c(\xi) - 2\bar{Y}(\xi)) + \frac{1}{2}\pi e \frac{\partial n}{\partial y}\Big(3(y - y_c(\xi))^2 + x^2\Big), \tag{4.53}$$

with the density at the geometric center of a given longitudinal slice given by

$$n_c(\xi) = n_0 + \left(\frac{\partial n}{\partial y}\right) y_c(\xi). \tag{4.54}$$

Next, we investigate the vertical gradient of the electron sheath's peak density. Here we combine the contributions of the electron sheath's density and axial current using n_{sh}. Figure 4.20(a)

shows that the peak of this combined electron sheath density and axial current have a linear dependence along the y dimension in the central region of the wake. This feature appears consistently in all the simulations in Sec. 4.3.1, except at the rear of the wake. Here we use a simple model of the electron sheath as a hollow tube of charge with a density profile given by

$$n_{sh}(x,y) = n_{0,sh} + \left(\frac{\partial n}{\partial y}\right)_{sh} y, \tag{4.55}$$

where $n_{0,sh}$ is the combined sheath density and current at $y = 0$ and $(\partial n/\partial y)_{sh}$ is the gradient of the combined sheath density and current. This sheath profile is defined over the region

$$R(\xi)^2 < x^2 + (y - y_c(\xi))^2 < (R(\xi) + L_{sh}(\xi))^2\,, \tag{4.56}$$

where $R(\xi)$ is the local wake radius and $L_{sh}(\xi)$ is the sheath thickness, which we approximate as

$$L_{sh}(\xi) \approx CR(\xi) \quad ; \quad C \equiv 0.0904 \tag{4.57}$$

with C as an empirical constant. The transverse wakefield within such a structure is fairly uniform, especially near the center, so we approximate the sheath's transverse wakefield everywhere as equivalent to the value at the center of this simplified charge distribution:

$$W_{y,sh}(\xi) \approx -2\pi e C R^2(\xi) \left(\left(\frac{\partial n}{\partial y}\right)_i - \left(\frac{\partial n}{\partial y}\right)_{sh}\right). \tag{4.58}$$

Here, $(\partial n/\partial y)_i$ is the density gradient of the ions from Eqn. 4.43. In addition to the longitudinally-varying wake radius $R(\xi)$ present in the other terms, Eqn. 4.58 relies on empirical models for L_{sh} and $(\partial n/\partial y)_{sh}$. In the simulations performed in Sec. 4.3.1, the sheath density gradient is fairly uniform over the bulk of the wake. We take the average linear fit of the sheath gradient over the longitudinal region $(1/2)\xi_{\max} < \xi < (3/2)\xi_{\max}$ and fit a linear relation using the ion density gradient to find

$$\left(\frac{\partial n}{\partial y}\right)_{sh} \approx B \left(\frac{\partial n}{\partial y}\right)_i \quad ; \quad B \equiv 3.19 \tag{4.59}$$

with empirical constant B.

We make the assumption here that the transverse wakefields within the blowout are entirely defined by the contributions from the ion column and from the combined sheath density and current

asymmetry. Therefore, after calculating the analytic wakefields from the ions using Eqn. 4.53, we assume the difference between Eqn. 4.53 and the wakefields in simulation is entirely from the sheath contribution of Eqn. 4.58. The empirical constants B and C are found first by taking the representative average of $(\partial n/\partial y)_{sh}$ and $W_{y,sh}/R(\xi)^2$ in the region $(1/2)\xi_{\max} < \xi < (3/2)\xi_{\max}$ for all simulations in Table 4.1. To account for numerical noise in the simulation, $(\partial n/\partial y)_{sh}$ and $W_{y,sh}/R(\xi)^2$ are smoothed using an 11-cell moving average algorithm. Figure 4.23 shows an example where the average $(\partial n/\partial y)_{sh}$ and $W_{y,sh}/R(\xi)^2$ were calculated for a simulation with $n_0 = 2 \times 10^{16}$ cm^{-3} and $g = 0.150$. It is worth noting in Fig. 4.23 that $W_{y,sh}/R(\xi)^2$ still has longitudinal variation. This suggests that the sheath model in Eqns. 4.55 and 4.57 is too simplified, and perhaps a more rigorous model that includes longitudinal evolution of the sheath's shape, gradient, and thickness can describe the wakefields more accurately. The averages of $(\partial n/\partial y)_{sh}$ and $W_{y,sh}/R(\xi)^2$ are then plotted in Fig. 4.24, where the empirical constants B and C are found by fitting Eqns. 4.59 and 4.58 across all simulations in Table 4.1 on a logarithmic scale.

Including the sheath contribution, the vertical wakefield becomes

$$W_y = 2\pi e n_0(\xi)\left(y - y_c(\xi) - 2\bar{Y}(\xi)\right) + \frac{1}{2}\pi e\frac{\partial n}{\partial y}\left(3\left(y - y_c(\xi)\right)^2 + x^2\right) + W_{y,sh}(\xi). \tag{4.60}$$

The total transverse wakefields from Eqns. 4.52 and 4.60 are plotted in Fig. 4.25 at the middle of the wake $\xi = \xi_{\max}$, where $R(\xi_{\max}) = R_{\max}$. The wakefields from the simulation in Fig. 4.25 are calculated using the total transverse electric, $\vec{E}_{\perp,\text{full}}$, and magnetic, $\vec{B}_{\perp,\text{full}}$, fields present in the simulation, $\vec{W}_\perp = \vec{E}_{\perp,\text{full}} - c\vec{B}_{\perp,\text{full}} \times \hat{\xi}$ [46]. The transverse wakefield equations include several terms with longitudinal variation and the implication of this variation is discussed in the following section.

Toward the rear of the wake, the sheath density gradient becomes much steeper than what the model predicts and, in cases with a large plasma density gradient g, the vertical sheath density profile becomes more quadratic or exponential. For such a density profile, the corresponding wakefield is no longer constant and a slight asymmetric focusing/defocusing perturbation develops. This, along with the limited region where the vertical offset drift is linear as seen in Fig. 4.21, may

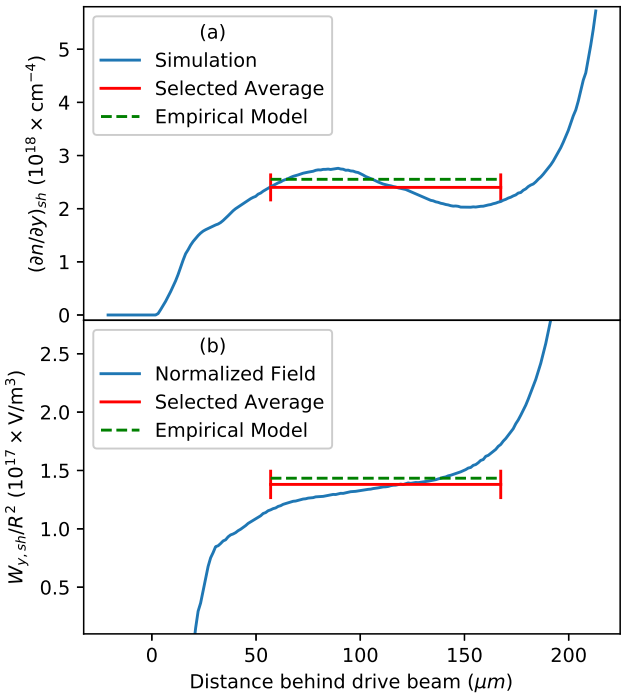

Figure 4.23: Sheath density gradient (a) and sheath wakefield (b) plotted (blue) along the longitudinal length of the plasma wake with $n_0 = 2 \times 10^{16}$ cm^{-3} and $g = 0.150$. The sheath density gradient and sheath wakefield curves are smoothed using an 11-cell moving average algorithm to suppress the effects of simulation noise. The horizontal red lines represent the average value in the interval $(1/2)\xi_{\max} < \xi < (3/2)\xi_{\max}$ and the vertical red lines are twice the standard deviation. The dashed green lines represent the empirical models of Eqns. 4.59 (a) and 4.58 (b), which utilize the mean of the selected average values (red lines) taken from all of the simulations. After 150 μm the model fails to accurately predict the density and field of the sheath.

add difficulty in applying this model to PWFA applications, where often the location of a witness beam for optimal acceleration and loading is in the rear of the wake.

4.3.3 Longitudinal Variation of Transverse Wakefields

To demonstrate the effectiveness of the preceding model, we consider the vertical deflection of a trailing witness bunch in an underdense, passive plasma lens, similar to that described in [66] but with a linear plasma density gradient. For such a lens, the witness bunch will have a longitudinally-dependent vertical deflection predicted by Eqn. 4.60. We perform another set of PIC simulations with both an electron drive bunch and a trailing witness bunch at various longitudinal separation

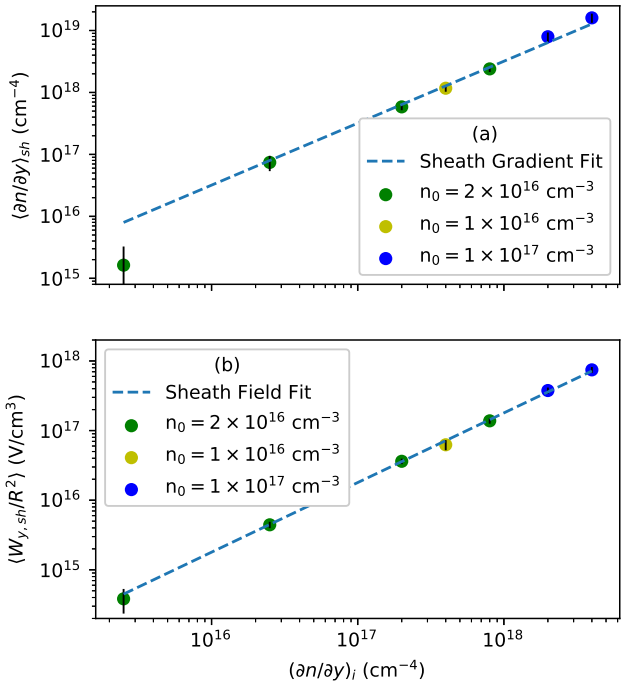

Figure 4.24: Comparison between the empirical models of Eqns. 4.59 and 4.58 (dashed blue) with the averages taken for the sheath density gradient (a) and sheath wake field (b) in the region $(1/2)\xi_{\max} < \xi < (3/2)\xi_{\max}$ across all simulation in Table 4.1 (dots). Errorbars are the standard deviation of these quantities within this region.

distances. The central plasma density and density gradient are set at $n_0 = 2 \times 10^{16}$ cm^{-3} and $\partial n/\partial y = 8 \times 10^{17}$ cm^{-4}, with $g = 0.150$. The normalized gradient used here is larger than for a typical gas jet profile, but it is useful for a more pronounced demonstration of this theory. Cases with a smaller normalized gradient will have a smaller effect on the deflection, but the dependence of the deflection angle on longitudinal position will be similar.

In addition to the inclusion of a witness bunch, a realistic model of the longitudinal plasma density profile is used, given by

$$n_p(x, y, \xi) = n_0 \left(\frac{1}{2} + \frac{1}{2}\tanh\left(\frac{\xi + a}{b}\right)\right)\left(\frac{1}{2} - \frac{1}{2}\tanh\left(\frac{\xi - a}{b}\right)\right)\left(1 + \frac{(\partial n/\partial y)}{n_0} y\right) \tag{4.61}$$

where $a = 150$ μm is the half-width, half-maximum (HWHM), and $b = 20$ μm determines the ramp steepness. This results in a plasma density profile with a flattop region of about 200 μm in length and density ramps with a full width of about 100 μm.

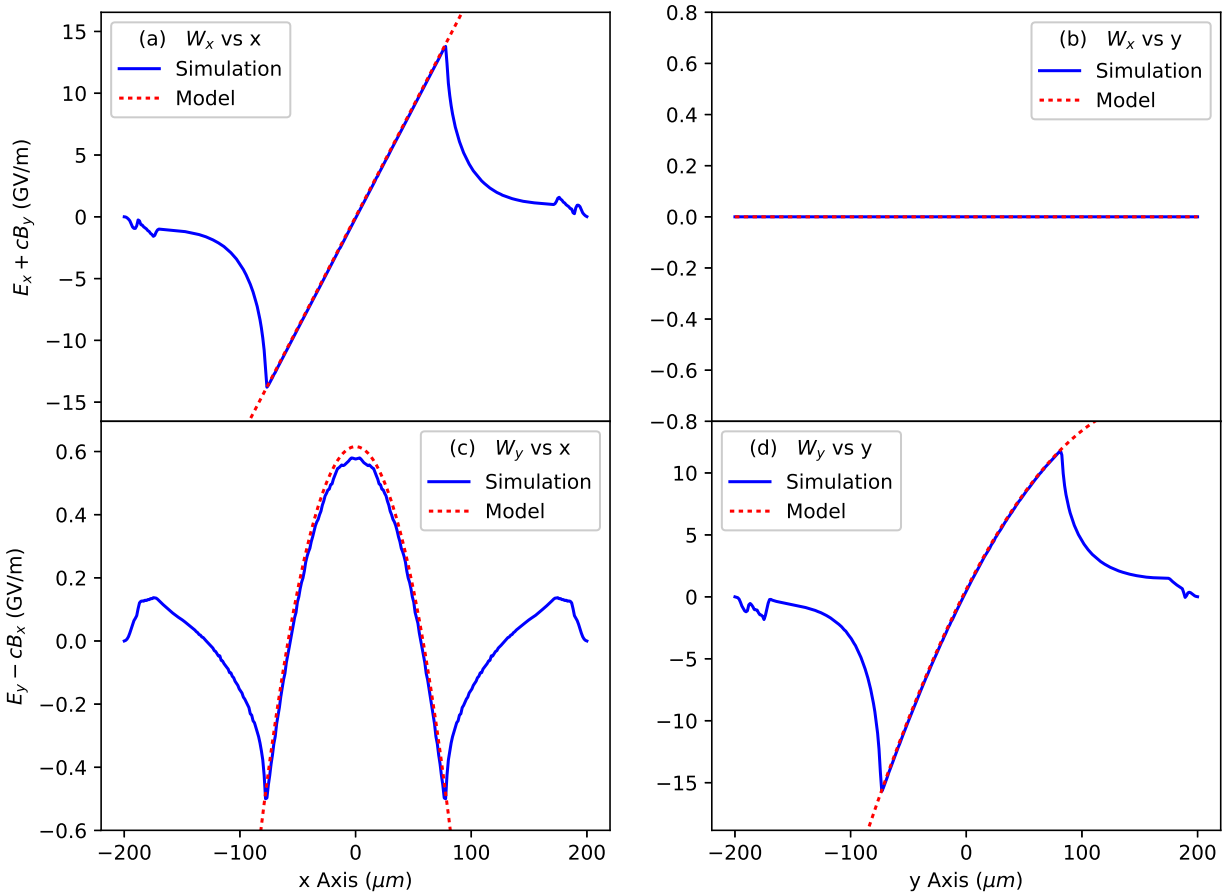

Figure 4.25: Transverse electromagnetic fields along the principle axes from a simulation with a central density $n_p = 2 \times 10^{16}$ cm^{-3} and a normalized transverse gradient $g = 0.150$. (solid blue), and from the model derived in this work (dashed red). Subplot (a) shows linear dependence on x for the horizontal focusing force; (b) shows zero variation of the horizontal focusing force along the y axis; (c) shows the vertical deflection and the small, nonlinear dependence on x of the vertical focusing force; and (d) shows the significant nonlinear dependence on y of the vertical focusing force.

The drive beam, plasma, and grid size parameters in these two-bunch simulations are all equivalent to those used in the corresponding single bunch simulations from Table 4.1. The witness beam parameters are as follows: energy $\gamma_L m_e c^2 = 10$ GeV, Gaussian energy spread $\sigma_\delta = 0.1$ %, longitudinal bunch length $\sigma_z = 6.0$ μm, transverse bunch size $\sigma_r = 6.0$ μm, normalized emittance $\epsilon_N = 3.1$ mm $-$ mrad, and charge $Q = 0.5$ nC.

The vertical deflection angle of the witness beam for each drive-witness separation distance is plotted in Fig. 4.26. We compare the deflection measured in the two-bunch PIC simulations

with the deflection predicted by Eqn. 4.60 using $R(\xi)$ and $\xi_{\max}$ from the zero gradient, steady-state simulation at $n_0 = 2 \times 10^{16}$ cm^{-3}. The deflection angle is calculated with respect to the drive beam axis in the longitudinal-vertical plane as

$$\theta_{\text{def}} = \frac{-eL_{\text{lens}}}{\gamma_L m_e c^2} W_y|_{y=0} \tag{4.62}$$

with e the elementary charge, m_e the electron mass, and $L_{\text{lens}} = 300$ μm the effective plasma lens thickness [66]. The deflection calculated from the electromagnetic fields of the single-bunch, long-propagation simulation from Sec. 4.3.1 with the same density and gradient is also plotted, assuming no density ramps. Here we consider this single-bunch simulation to be effectively steady-state, as the wake has stopped evolving after 5 mm of propagation distance. The agreement between these plots is good throughout most of the wake, but we can make a few important observations from the discrepancies between them.

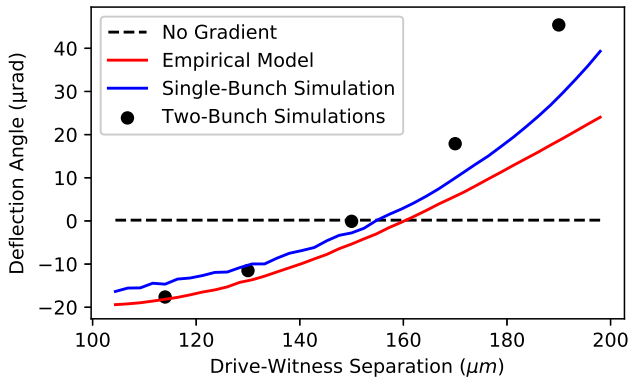

Figure 4.26: The longitudinal variation of the net deflection angle (left axis) from an underdense plasma lens with FWHM = 300 μm, $n_0 = 2 \times 10^{16}$ cm^{-3}, and $g = 0.150$ for a witness beam behind and on the same axis of a drive beam. Solid red is the prediction from the empirical model in Eqn. 4.60, solid blue is the prediction from the steady-state simulation of the same gradient and not including density ramps, and the black dots are taken from two-bunch simulations in a short plasma lens density profile. The maximum blowout radius in these simulations occurs at a distance $\xi_{\max} = 112.2$ μm behind the drive beam.

First, the difference between the empirical model and the steady state fields arises from the breakdown of the assumptions in the empirical model toward the rear of the wake. With the high density sheath approaching the drive beam axis more rapidly than the low density sheath, the

blowout wake cross section begins to resemble more of a heart shape than a circle. We can interpret from Fig. 4.21 that this departure from a roughly circular cross-section takes place beyond 200 μm behind the drive beam, at which point the model will no longer accurately reflect the fields of the ion column. However, in Fig. 4.26 the results diverge well before 200 μm, indicating this is not the primary source of error. Rather, from Fig. 4.23 we can see that the model for the sheath's wakefield is only valid in the center of the wake and diverges around $\xi = 150$ μm because the density gradient of the electron sheath becomes much steeper than the model predicts.

Second, the difference between the single-bunch steady state fields and the results of the two-bunch simulations is due to the longitudinal density profile used in the latter. With this density profile, the beams do not propagate long enough through the flattop region of the plasma to reach a steady state plasma wake. The density ramps also contribute to these non-steady state conditions. We see significant variation in the instantaneous transverse wakefield that appears to oscillate about the expected steady-state transverse wakefield for the short duration of the plasma lens. Furthermore, it is possible that the transverse wakefield of the witness beam itself has enough of an effect on the sheath of the blowout that the rear of the witness beam is affected.

We here make a general observation from Eqn. 4.60 that a witness beam located toward the middle or front of the wake will tend to experience an ion-dominated deflection toward the direction of the gas jet nozzle, while a witness beam near the rear of the wake will experience a sheath-dominated deflection away from the gas jet. This trend is clearly observable in Fig. 4.26.

4.3.4 Transverse Variation of Longitudinal Wakefields

Since there is a longitudinal variation of the transverse focusing fields in the blowout wake, the Panofsky-Wenzel theorem predicts a corresponding transverse variation of the longitudinal accelerating fields [91, 92]. This is indeed what we observe in simulations as shown in Fig, 4.27, and we can calculate the slope of W_z using Eqns. 4.52 and 4.60 and solving

$$\frac{\partial W_z}{\partial x} = -\frac{\partial W_x}{\partial \xi}; \qquad \frac{\partial W_z}{\partial y} = -\frac{\partial W_y}{\partial \xi}. \tag{4.63}$$

We find the following expression for the longitudinal wakefield:

$$W_z = \frac{1}{2}\pi e\alpha \frac{\partial n}{\partial y}(x^2 - y^2) + \Omega_{\rm lin} y + W_{z,0}(\xi), \tag{4.64}$$

where α is defined from Eqn. 4.50 and $W_{z,0}(\xi)$ is the longitudinal wakefield along the drive beam axis, and $\Omega_{\rm lin}$ is defined

$$\begin{aligned}\Omega_{\rm lin} \equiv \pi e \Bigg[& \alpha^2 \frac{\partial n}{\partial y}\xi - 2n_0\alpha - \alpha\left(\frac{\partial n}{\partial y}\right)^2 \frac{R^2(\xi)}{n_0} \\ & + 2R(\xi)\frac{\partial n}{\partial y}\frac{\partial R(\xi)}{\partial \xi}\left(1 - \frac{\alpha}{n_0}\frac{\partial n}{\partial y}\xi\right) \\ & + 4CR(\xi)(1-B)\frac{\partial n}{\partial y}\frac{\partial R(\xi)}{\partial \xi}\Bigg].\end{aligned} \tag{4.65}$$

The main contributions to a non-flat longitudinal wakefield are found in Eqn. 4.65. In particular, the largest terms are the quantities $2n_0\alpha$ and $2R(\xi)(\partial n/\partial y)(\partial R(\xi)/\partial \xi)$. If we consider the middle of the wake ($\xi = \xi_{\rm max}$), the longitudinal evolution of the wake radius drops out of Eqn. 4.65 and we find a much simpler expression:

$$\Omega_{\rm lin,mid} = \pi e\alpha \left[\alpha \frac{\partial n}{\partial y}\xi_{\rm max} - 2n_0 - \left(\frac{\partial n}{\partial y}\right)^2 \frac{R_{\rm max}^2}{n_0}\right], \tag{4.66}$$

which for realistic parameters can be approximated

$$\Omega_{\rm lin,mid} \approx -2\pi e\alpha n_0. \tag{4.67}$$

Equation 4.64 agrees well with observations in the front of the wake. As with the vertical transverse wakefield, the longitudinal wakefield plots in Fig. 4.27 agree less with Eqn. 4.64 towards the rear of the wake. We also note that the position in the wake that corresponds to a flat longitudinal wakefield is in the front half of the wake between 60 and 90 μm, approximately where the slope of the transverse wakefield would be zero in Fig. 4.26.

There is no point along the drive beam axis where both the transverse field is zero and the longitudinal field is uniform. However, these conditions can be satisfied at a transverse offset position in the wake. In this case, one would solve Eqn. 4.65 to find ξ where the longitudinal wakefield slope is zero, then solve Eqn. 4.60 to find the needed vertical offset from the drive beam

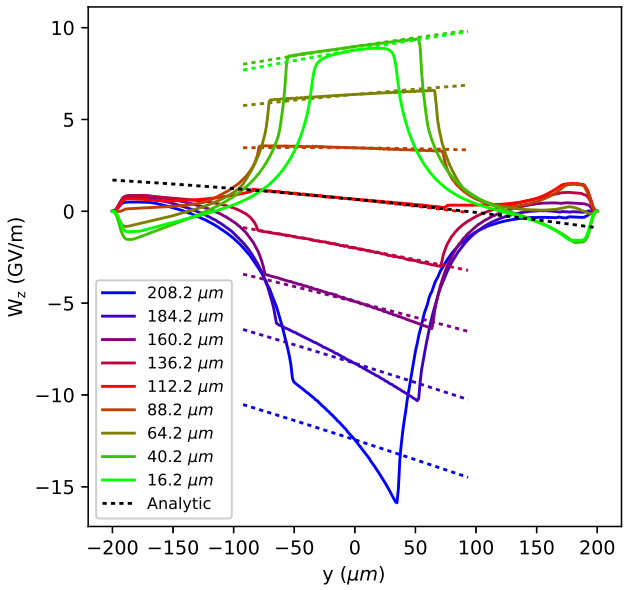

Figure 4.27: Transverse variation of the longitudinal wakefield at various longitudinal positions in the blowout wake behind the drive beam. For example, the light green curve is 16.2 μm behind of the center of the drive beam, and the dark blue curve is 208.2 μm behind the drive beam, close to the rear of the wake. Dashed lines of the corresponding color represent the calculated slope using Eqns. 4.64 and 4.65. The slope of this curve at the center of the wake ($\xi = 112.2$ μm) matches the analytic solution of Eqn. 4.66 (dashed black).

axis to achieve zero deflection. This solution would be in the front half of the wake with a slight decelerating field, sufficient for a plasma lens. But in the context of a plasma accelerator that requires an accelerating field, it would be impossible to find a longitudinal position inside the wake with a transversely uniform accelerating field.

4.3.5 Implications of a Linear Density Gradient Plasma

The linear density gradient is a good approximation for the plasma profile produced in a gas jet outflow, as the length scales of the plasma wake are typically small compared to the characteristic length scales of the gas outflow. The variation of the fields in the plasma wake with respect to the uniform density case arise from the transverse density gradient of the ions and the sheath electrons, as well as the longitudinally drifting transverse center of the wake. This drifting of the wake's center is similar to what was seen when using an elongated, tilted drive beam [93]. A theoretical model of

the field structure was derived based on an analytic approach supplemented with empirical models based on observations from PIC simulations.

The model presented here does not have an unlimited range of applicability. There are two constraints in particular that one should consider. First, the density gradient must be linear or very close to linear. This condition will likely be satisfied in the outflow of a typical gas jet so long as the blowout radius is small compared to the characteristic length scale of the gas jet's exponential density profile. Second, the plasma density must not fall to zero within the blowout radius, otherwise the wake would deform into a non-circular shape early in the blowout. This implies that the normalized density gradient radius must satisfy the condition $gR_{\max}\omega_{pe}/c < 0.5$.

We applied this model to predict the net deflection of a witness beam in a two-bunch simulation propagating through a 400 μm thick plasma lens with a realistic density profile and an exaggerated density gradient compared to a realistic gas jet. The model worked well when the witness beam was located within the bulk of the plasma wake, diverging from the simulation results only when approaching the rear of the wake where our assumption of a circular cross-section and a constant, linear sheath density gradient break down. Within the region $(1/2)\xi_{\max} < \xi < (3/2)\xi_{\max}$ where the electron sheath model was empirically obtained, the model agrees very well. Additional factors that limited the accuracy of this model were the density ramps and an overall length that was too short to allow a steady state to be reached.

The net deflection of the witness beam depends on its longitudinal location within the wake, and shifts from an ion-driven deflection toward the higher density side when near the front of the wake, to a sheath-driven deflection toward the lower density side when near the rear of the wake. This model allows for the prediction of the correct position of the witness beam to prevent a net transverse deflection despite the presence of the transverse density gradient. Alternatively, one could design a plasma lens with a longitudinally varying central density n_0 that results in non-zero deflections throughout the plasma but a net-zero deflection overall.

In addition, this model describes the transverse dependence of the longitudinal wakefield in the wake. For a thin plasma lens, this is likely to be a negligible effect. However, this model applies

equally well to a plasma wakefield accelerator operating in the nonlinear blowout regime, and in this context the longitudinal field variation could be of great consequence. This model can predict the magnitude of this effect, which may be of critical importance when designing a plasma-based accelerator in an elongated gas jet that aims to preserve beam quality (e.g. emittance, energy spread), such as presented in Ref. [94, 95].

While the primary focus of this study is electron beam-driven nonlinear blowout wakes, many of the same conclusions can be drawn for a high-intensity laser-driven wakes. An important caveat, however, is that the front of the laser driver will experience a transverse deflection due to the transverse gradient of the plasma refractive index [96]. Meanwhile, the behavior of an electron beam driver would be somewhat more complex. the tail of the beam inside the blowout would oscillate about the center of the ion column's focusing force. This may, in turn, lead to an average drift of the wake center over time as the head of the beam erodes. The dynamics of the head of the beam, however, are difficult to model and predict with certainty without parameter-specific simulations.

It is worth noting that the effects of the transverse density gradient on the performance of a thin, underdense plasma lens may not always be of negative value. A particular asymmetric density profile can be designed as to make a plasma kicker, similar to what is discussed in Refs. [97, 98]. A plasma kicker with a large normalized density gradient might find use as a compact replacement for a dipole magnet when integrating or separating drive and witness bunches in a multi-stage, plasma-based linear accelerator [99]. Under the correct conditions, it could even potentially be used as a compact alternative to a sextupole magnet for correcting beam aberrations with a more forgiving beam alignment tolerance. One could even design a plasma-based beam streaker that is capable of deflecting the head and tail of an electron beam at different angles, as evidenced by Fig. 4.26.

In the context of the passive plasma lenses relevant for this thesis, this effect is something that should eventually be studied in experiment. Typical gas jet outflows will have a fairly small gradient where the density varies by a few percentage over the width of the blowout, but this

can still result in sub-optimal performances if a plasma lens is designed to focus beams with very accurate precision. While we will use a gas jet for the plasma lens design discussed in Chapter 5, testing the effects of the nonuniform density profile is beyond the scope of the experiment in this current, early phase.

Chapter 5

Experimental Setup

In Chapter 3 we considered plasma lenses from a theoretical perspective and develop an understanding of their capabilities and requirements. Then, Chapter 4 explored methods of laser ionizing plasma lenses with acceptable parameters using feasible laser and gas systems. With these design considerations in mind, we now move to the experimental efforts to realize these laser-ionized plasma lenses with an electron beamline. In this Chapter, we introduce the experimental facility, FACET-II, that these plasma lens experiments will be carried out in. This includes a 1 km electron accelerator (Sec. 5.1), a 10 TW laser system (Sec. 5.2), and an accompanying suite of diagnostics in the experimental area for measuring the laser, plasma, and electron beam (Sec. 5.3).

As of the time of writing, FACET-II is still in the later stages of commissioning. The laser system is being continually improved, and the electron beam parameters are still in the process of being tuned to their final goals. For the time being, we only have access to a single electron bunch rather than the two-bunch configurations for ideal plasma lens operation. We also were unable to generate a plasma using the low energy probe with the spherical lens focused discussed in Sec. 4.1.3, and had to instead pivot towards a longer, more reliably ionized plasma with a separate laser line. Regardless, the work we have performed on the experimental setup for the plasma lens experiment will be useful in the future once the laser and electron beam capabilities have been improved.

5.1 FACET-II

The electron beam source we will be using is that produced by the 1 km long accelerator at the Facility for Advanced Accelerator Experimental Tests (FACET-II), located at SLAC National Accelerator Laboratory in California. FACET-II is the successor facility to FACET, which ran from 2011 to 2016 and occupied the first 2 km of SLAC's full 3 km accelerator [8]. During this time, FACET delivered 20 GeV electron and positron beams to experiments located at the end of this linac section. In 2016, the first km of FACET was redesignated for the LCLS-II accelerator and FACET was left with only the middle km of the SLAC linac remaining. Parallel to this change FACET began upgrading towards FACET-II using a new photo-injector electron beam source. So even though the maximum electron beam energy is halved to 10 GeV, FACET-II is designed to enable an electron beam emittance that is significantly lower and a peak current that is much higher. This redesign of the accelerator and experimental area enables many research experiments in high energy physics and advanced acceleration techniques [11].

In this section, we present a summary of FACET-II, and give the electron beam parameters during the time of commissioning and the anticipated final parameters after installation and commissioning is completed.

5.1.1 FACET-II Linac

SLAC's full 3 km linac is divided into 30 sectors, each a tenth of a km long. FACET-II begins in sector 10 with the electron gun and ends with the experimental area in sector 20, followed by a lead brick shielding wall. The electron accelerator itself is located about 10 m underground, and this tunnel also hosts all of the magnets and RF cavities used to control and accelerator the beams. Above the tunnel on the ground floor is the Klystron gallery, which houses the RF amplifiers (Klystrons) that power the RF accelerator cavities. The gallery also contains numerous electric and vacuum equipment that would be too susceptible to radiation damage if it were located in the accelerator tunnel. At the experimental area, the Klystron gallery also contains all of the

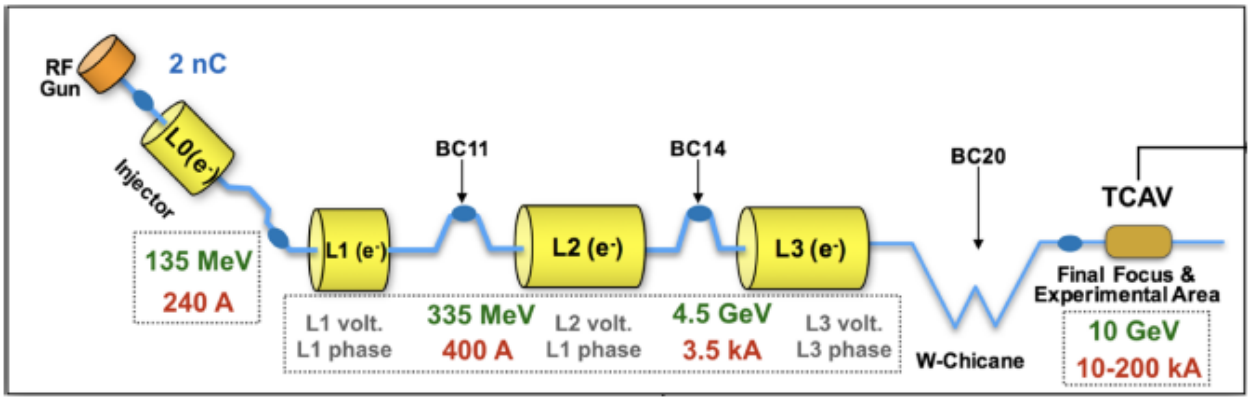

Figure 5.1: Schematic of the FACET-II accelerator. The electron beam source is the RF gun representing the photo-injector. The electron then travels through several RF accelerating structures (yellow) and bunch compressors (BC) before arriving at the experimental area. Reproduced from Ref. [100]

electronic hardware used by the experiments, such as camera servers and oscilloscopes, with long cables connecting the servers to the equipment in the accelerator tunnel.[1]

A top-level schematic of the FACET-II accelerator is shown in Fig. 5.1. The electron injector at Sector 10 is a photocathode source. This operates by sending a 400 μJ, $\lambda = 253$ nm UV laser pulse onto a copper cathode, and these high energy photons then eject 2 nC of electrons into the first RF cavity. The RF phase is synchronized with the arrival of the UV laser to be able to accelerate the ejected electrons, and this new beam is quickly accelerated to 4.2 MeV before the emittance grows from space charge effects. A solenoid then focuses this beam to control its divergence. A drawing of this electron source and subsequent structures are shown in Fig. 5.2. Afterwards, the beam travels through two 3 m S-band accelerating structures that accelerate the beam to 135 MeV. Taken together, this represents L0 in Fig. 5.1.

After L0, the electron beam is sent onto the accelerator's main axis by dipole magnets. For the next 1 km until the experimental area at sector 20, the beam travels through many accelerating structures and bunch compressors. The bunch compressors, or chicanes, are comprised of four or more dipole magnets that use the beam's energy spread to reduce its longitudinal size. As shown

[1] Cameras in the tunnel do feel the effects of radiation over time from being so close to the electron beam. This is the inspiration for FACET-II's unofficial slogan: "Where cameras go to die"

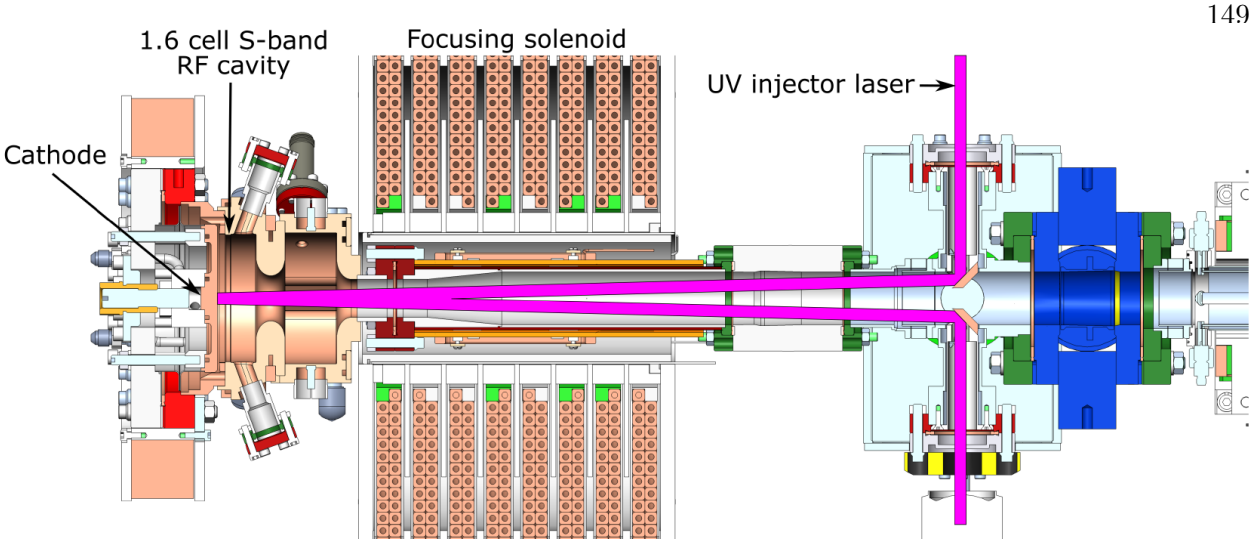

Figure 5.2: The photo-injector electron source at FACET-II. The UV laser pulse reflects off of a copper cathode which ejects 2nC of electrons into the RF accelerating cavity and subsequent focusing solenoid. After this injector, the electron beam reaches an energy of 4.2 MeV. Reproduced from Ref. [33]

in Fig. 5.1, the accelerating structures are grouped into three stages, with each stage consisting of many 3 m long S-band RF cavities.[2] In acceleration stage L1, the beam is accelerated to 335 MeV with a $-20.5°$ phase offset between the RF pulse and the electron beam's arrival. This phase offset generates an energy chirp, which is used by the subsequent bunch compressor (BC11 in sector 11) to reduce the beam's longitudinal size to 468 μm through four bends with dipole magnets. In accelerating stage L2, the beam accelerates to 4.5 GeV with a $-39.0°$ phase offset, and is compressed with four bends at BC14 to 96 μm bunch length. Here, the beam is over-compressed so that it is chirped for the final compressor. After L3, the beam is at 10 GeV with no phase offset. The final "W-chicane" bunch compressor compresses the length to 1.8 μm.

This last bunch compressor and subsequent experimental area is shown in the diagram of Fig. 5.3, which represented the hardware in place during FACET. Most notably, final focus magnets bring the electron beam to its smallest focus at the experimental area (or IP, for "interaction point"). After the IP, several magnets can image the electron beam onto downstream diagnostic screens. This imaging spectrometer is discussed further in Sec. 5.3.3.

[2] A single Klystron can power up to three of these 3 m RF stages.

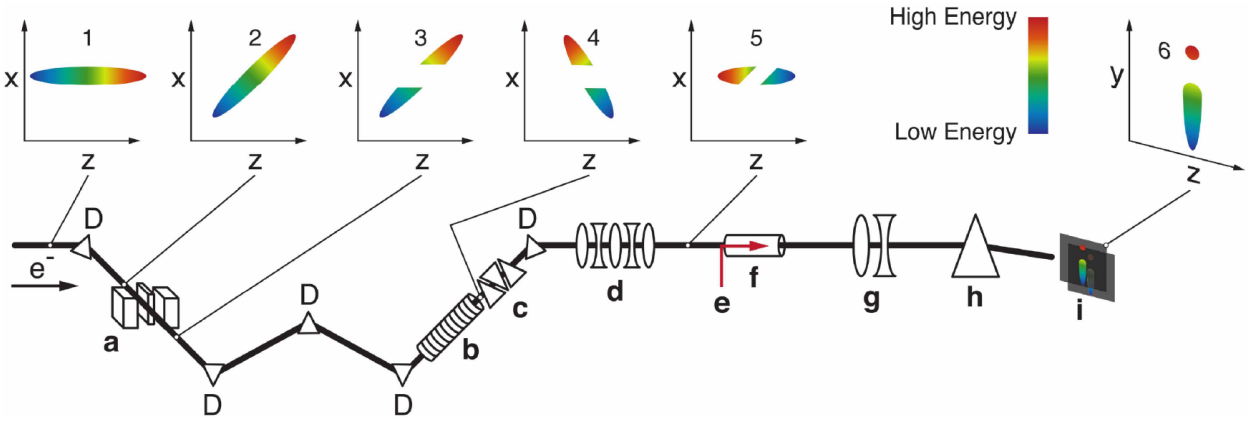

Figure 5.3: Diagram of the W-chicane bunch compressor and experimental area of FACET. Dipole magnets are labeled with a D, (a) is the notch collimator, (b) is the transverse deflecting cavity (TCAV), (c) is the a spectrometer, (d) is the final focus magnets, (e) is the ionization laser, (f) is the plasma, (g) is the imaging spectrometer quadrupole magnets, (h) is the dipole magnet for the imaging spectrometer, and (i) are screens for imaging the electron beam. At FACET-II, the TCAV is moved after the chicane and an additional quadrupole magnet is installed in the imaging spectrometer. Reproduced from Ref. [9]

Lastly, we mention here how to generate a two-bunch electron beam structure. Up to this point, we have only been considering the FACET-II accelerator in the context of accelerating a single electron bunch. There are two methods to achieving a two-bunch structure at FACET-II. The first, also presented in Fig. 5.3, is to use the energy chirp present in the single electron beam upstream of the W-chicane. After the beam is dispersed transverse following the first bend in the chicane, a thin, dense metal (called a notch collimator) can be inserted to block out a portion of the beam. Once this beam is recompressed in the chicane, the result is a beam divided into two bunches with a longitudinal separation. The second method involves the photo-injector at the beginning, where we can instead send two UV pulses onto the copper cathode and generate two electron bunches with a longitudinal separation at the start of the linac. While the latter method would be more desirable for beam quality, it is more difficult to optimize two electron beams through the full linac rather than using the quick and messy approach of the notch collimator.

5.1.2 Parameter Space

The nominal electron beam parameters are obtained by running particle tracking simulations for the entire accelerator and optimizing electron beam parameters for various experiments. This is what the FACET-II team at SLAC has done for a two-bunch configuration with a drive and witness bunch, and by taking the 6D phase space distributions at the end of these simulations we can generalize the electron beam distributions with Gaussian fits and CS parameter definitions from Chapter 2. The final focusing quadrupole magnets just upstream of the experimental area are capable of delivering a range of parameters around these nominal values, where it is generally easier to request larger beams than it is to request smaller beams.

With the witness bunch's vacuum waist at the IP, the CS parameters for the witness bunch are $\alpha_{x,y} = 0.1, 0.7$ and $\beta_{x,y} = 5.0, 7.0$ cm. At this longitudinal position, the leading drive bunch is upstream of its focus (the horizontal x focus is 14.8 cm downstream and the vertical y focus is 34.1 cm downstream) and so the drive bunch is much larger with $\alpha_{x,y} = 1.7, 2.2$ and $\beta_{x,y} = 33, 33$ cm. This difference in focus location is due to the difference in beam energy between the drive and witness bunches, which shows up as chromatic aberrations in the strong, final focusing magnets. The emittances for these electron beams in the two-bunch configuration are designed to be $\gamma\epsilon_{x,y} = 3.7,\ 4.1\ \mu\text{m} - \text{rad}$ for the witness bunch and $\gamma\epsilon_{x,y} = 35.0,\ 115\ \mu\text{m} - \text{rad}$ for the drive bunch. The repetition rate for these beams is nominally set at 10 Hz.[3]

The above parameters are the designed FACET-II electron beam parameters for a two-bunch beam that is completely tuned up. At the time of writing this dissertation, FACET-II is still finishing its commissioning phase and the accelerator is running with only a single electron bunch. The parameters for this single bunch are being tuned up as well, and the electron beam that was available to us during the plasma lens commissioning shift will be discussed later in Chapter 6. That being said, with access to a single 10 GeV electron bunch with high peak current there is still plenty of science that can be accomplished in parallel to the accelerator commissioning progress,

[3] It is easy to get lower repetition rates by blocking the electron beam on particular shots, but going to higher repetition rates is limited by the UV laser at the photo-injector.

including early plasma lens experiments.

5.2 Laser

Along with the electron beam itself, a critical component of the plasma lens experiment is the ionizing laser. In Chapter 4 we discussed the necessary laser parameters in order to ionize a plasma lens of the desired shape. In this section we introduce the layout of the FACET-II laser system at the experimental area of sector 20. We also discuss two available methods for ionizing a plasma lens, and future work to improve on these laser ionization schemes.

The FACET-II laser system is a 10 TW ultrashort pulsed laser operating in the IR range at $\lambda = 800$ nm. The current laser system is very similar to that which was available during the FACET years [101]. It is a chirped-pulse amplification (CPA) system, with multiple amplification stages all located in an initial laser room on the ground floor. Once the final laser energy and profile is reached, the uncompressed laser pulse is sent downwards through the laser transport and into the experimental area in the accelerator tunnel. Once in the tunnel the laser is split into a "main" line and a "probe" line with a beam splitter, each line gets compressed and sent to experiments.

5.2.1 Sector 20 Laser Room

The sector 20 (S20) laser room is where many key components of the CPA laser system are located (Fig. 5.4). The laser starts at an oscillator (Coherent Inc. Vitara-T), which produces ultrashort pulses at 800 nm wavelength, 60 nm FWHM bandwidth, and a repetition rate of 68 MHz. This pulse train output from the oscillator is phase-locked to the RF signal from the linac, and all laser components are timed to this oscillator output so that the laser can be synchronized with the electron beam's arrival. However, there does exist jitter between the laser and electron beam of about 48 fs due to jitter in the time of arrival for both.

After the oscillator, the pulses are sent into a regenerative amplifier (regen, Coherent Inc. Legend Elite HE USP). Here the pulses are stretched and subsequently amplified to 1 mJ with a repetition rate of 120 Hz. Afterwards, a pulse picker reduces the repetition rate down further to

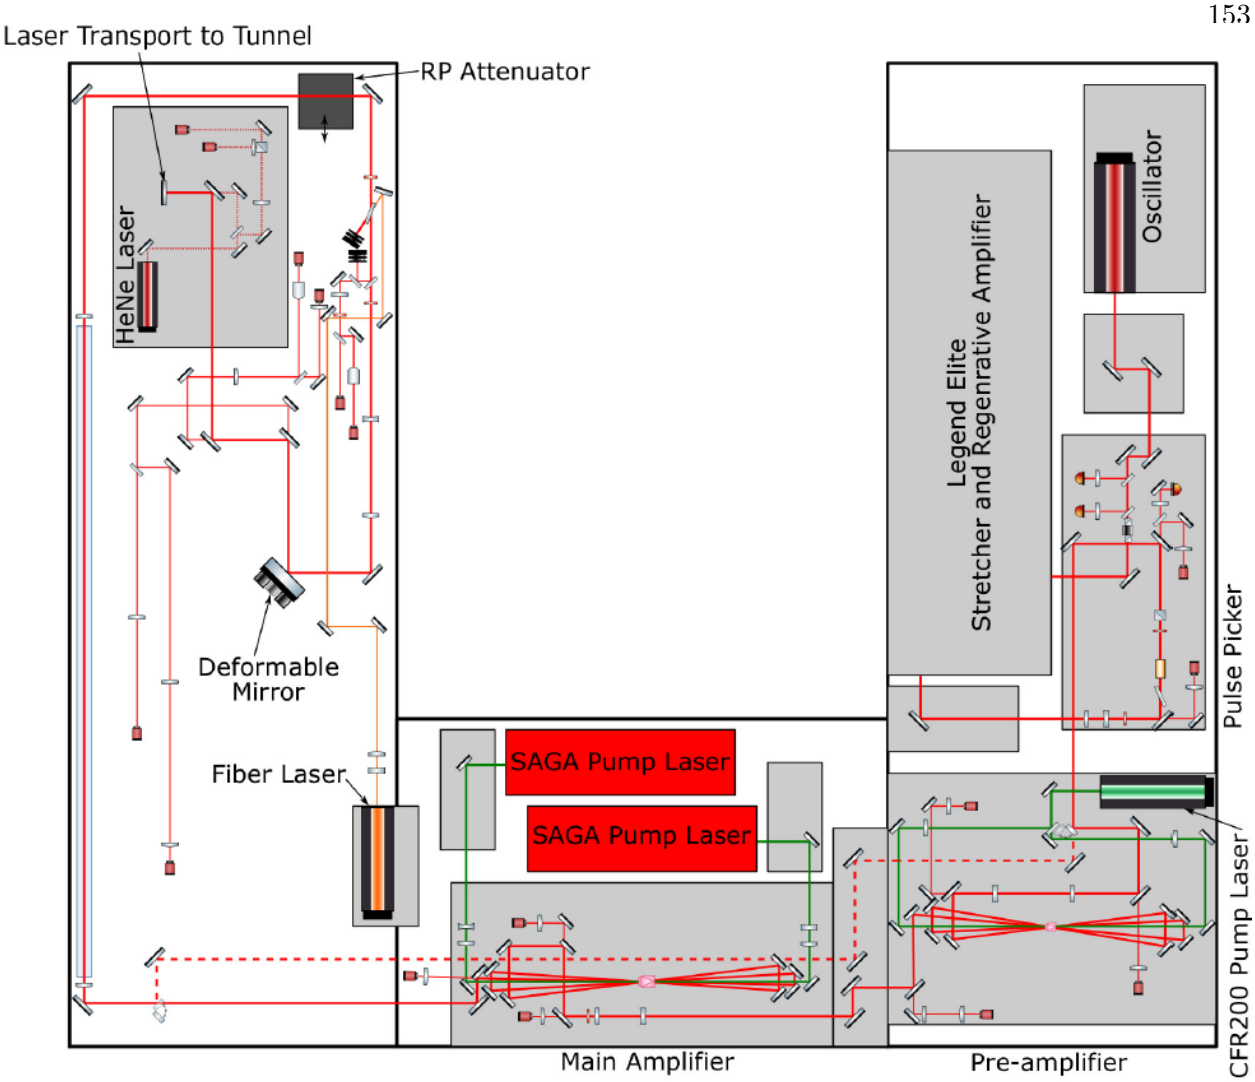

Figure 5.4: Overview of the FACET-II laser system in the sector 20 laser room. The laser starts at the commercial oscillator in the top-right before passing through several amplifier stages. After the main amplifier, the laser passes through a relay imaging telescope, energy attenuators, a beam expanding telescope, and a deformable mirror before being sent into the transport towards the tunnel. Reproduced from Ref. [33]

either 1 or 10 Hz depending on the experimental requirements.

The next stage is a four-pass preamplifier that brings the laser pulse energy up to 30 mJ. This preamplifier is pumped by a 130 mJ Nd:YAG laser (Quantel CFR200) that is frequency doubled to 532 nm. After the preamplifier is the main amplifier that boosts the laser energy to 1 J. The main amplifier is pumped with a pair of 1.8 J, 532 nm frequency doubled Nd:YAG lasers (Thales

SAGA) onto a Ti:sapph crystal. After amplification, the bandwidth is 24 nm FWHM and the laser diameter is about 10 mm. Here the laser energy is maximized at 800 mJ per pulse.

This laser pulse is then sent through a suite of diagnostic cameras, safety systems,[4] and beam shaping optics (Fig. 5.5). A 4f image relay telescope preserves the laser profile immediately after the main amplifier. An energy attenuator consisting of a motorized waveplate and a polarizer allow for the energy sent into the tunnel to be attenuated. Next, a diverging-converging telescope expands the beam from 10 to 40 mm. After the beam is expanded to 40 mm, it also reflects off of a deformable mirror (DM). The DM has the capability to flatten the wavefront of the laser pulse by using actuators within the mirror to deform the surface of the DM and correct for any phase aberrations in the laser pulse. This is done by pairing the DM with a Shack-Hartmann wavefront sensor that images the surface of the DM using a 4f imaging system. A closed loop algorithm between the DM and wavefront sensor then produces the desired wavefront in the laser pulse.

All throughout this laser system are cameras and motorized mirrors that allow for the remote alignment of the laser system. An auto-alignment system is in place to adjust the steering of mirrors so that the laser pulse is hitting target positions on cameras along the laser path. This system is a closed-loop algorithm that continually runs while the laser is turned on. This is especially important to compensate for long term drift, which arises due to the laser room being on the ground floor and susceptible to outside variations in temperature and weather.[5]

5.2.2 Laser Transport

The "laser transport" refers to the section of the laser system that connects the Sector 20 laser room with the experimental area underground. This transport consists of a 30.5 m long in-vacuum tube with seven mirrors (these are called B0 - B6) before entering the main laser compressor chamber in the tunnel. The transport begins with the "launch mirror" (labeled "M7" in Fig. 5.5) located below a vacuum window flange. After the launch mirror, the laser propagates upwards to

[4] By design, the FACET laser can only be sent into the tunnel at full power when there are no personnel in the tunnel. This is the case when the electron beam is running. During alignment in the tunnel, the laser energy is reduced to 1% energy with an insertable radiation protection (RP) attenuator.

[5] And the occasional vehicle.

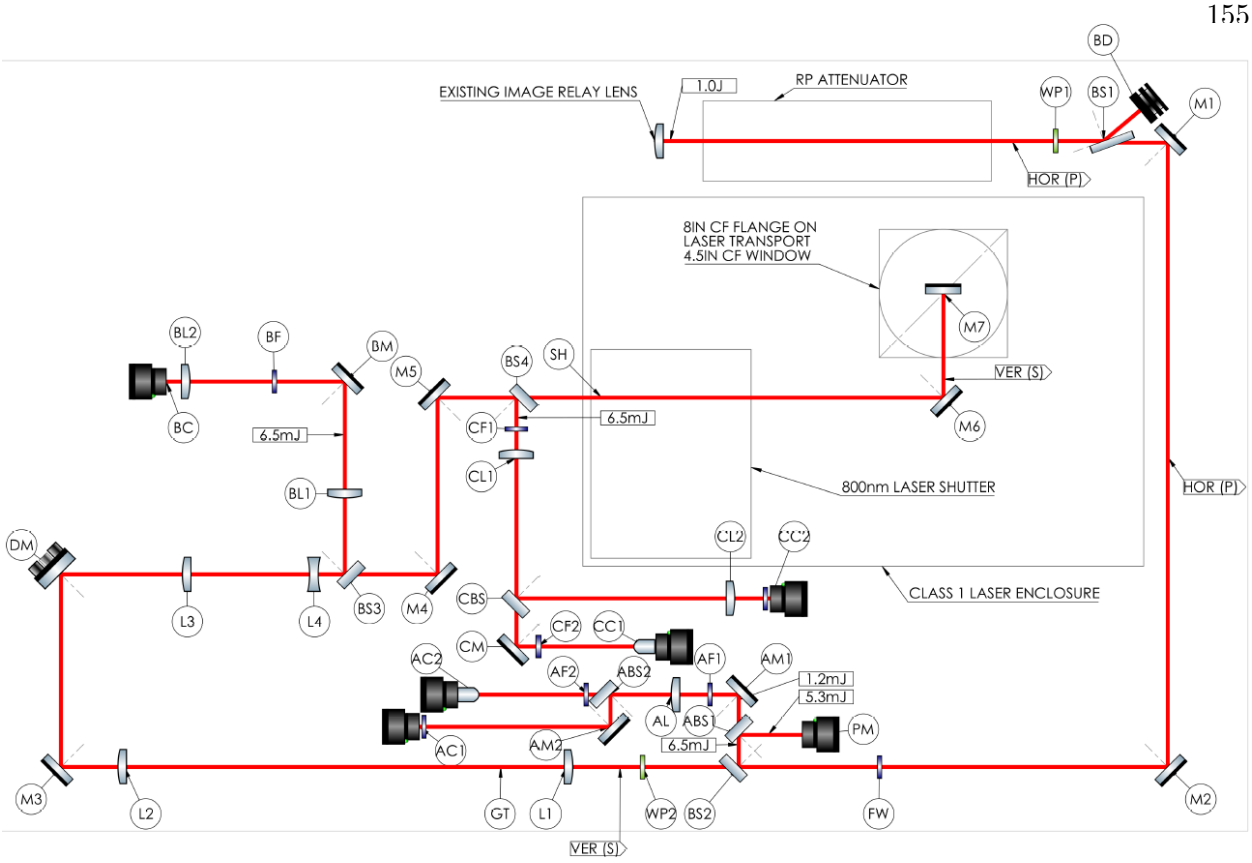

Figure 5.5: Schematic of the laser system between the main amplifier and the tunnel transport. The laser enters here at the top "existing image relay lens" and travels through the radiation protection (RP) attenuator, another independent energy attenuator, several diagnostic imaging systems that use leakage light from mirrors, and a deformable mirror. A laser shutter is in place before the transport to allow for the laser to be blocked remotely. Reproduced from Ref. [33]

mirror B0, then exits the laser room towards the Klystron gallery and mirror B1. The laser travels along the Klystron gallery to mirrors B2 and B3, and at B3 the laser is sent directly downwards 11 m towards the experimental area in the accelerator tunnel. Once in the tunnel, the laser first reflects off of beam splitter B4 and travels upstream in the tunnel to mirrors B5 and B6, which reflect the laser down into the large compressor chamber adjacent to the vacuum chamber at the IP used for experiments. Discussion of the main laser line is continued in Sec. 5.2.5. The low energy probe laser line is extracted through the transmitted path below B4 and contains 20% of the full laser energy. The probe line exits the vacuum transport system through a window and continues downwards towards a separate pulse compressor. Discussion of the probe laser is continued in

Sec. 5.2.3.

All of the mirrors in B0 through B6 have cameras that use the leakage light to track the position of the laser in the transport. The exception is the beam splitter B4, but there is a camera that looks at leakage light in the first mirror of the probeline directly after the B4 leakage. To assist with alignment through this laser transport, a HeNe laser is injected co-linear with the main beam at the initial launch mirror. This HeNe laser is smaller and more accurate to align than the large 800 nm beam, so it is left on while an autoaligner system continuously corrects for long-term drifts in the transport using the HeNe. This is important since the laser transport exists in both the temperature-controlled sector 20 laser room, the ground floor Klystron gallery, and the underground accelerator tunnel. The Klystron gallery is not temperature-controlled and components of the transport that are mounted to the Klystron gallery walls are susceptible to fluctuations from windy weather. Additionally, conditions in the tunnel depend on the state of the accelerator. When the accelerator is running, the 80° F temperature set-point for the magnets' chilling water can increase the tunnel's temperature significantly and further make laser alignment difficult. The HeNe autoalignment and camera references thus are crucial in providing an easy way to manage these conditions.

5.2.3 Laser Probeline

The laser probeline begins in the experimental area below the B4 beam splitter in the transport vacuum tube. Here 20% of the laser energy is transmitted to the probeline and the laser is still uncompressed. Figure 5.6 is a drawing of this initial section of the probeline. After B4, the laser is sent downwards to the table height and then over towards two steering mirrors that align into a reducing telescope that reduces the beam diameter to 10 mm. Then a waveplate/polarizer pair acts as an energy attenuator to further limit the laser energy if needed. The probe is then steered into a compressor, which compresses the laser pulse to a pulse duration less than 70 fs. The beam exits this compressor box at a lower beam height, where two mirrors subsequently steer the probe upstream in the accelerator tunnel along the side of the experimental vacuum chamber. In

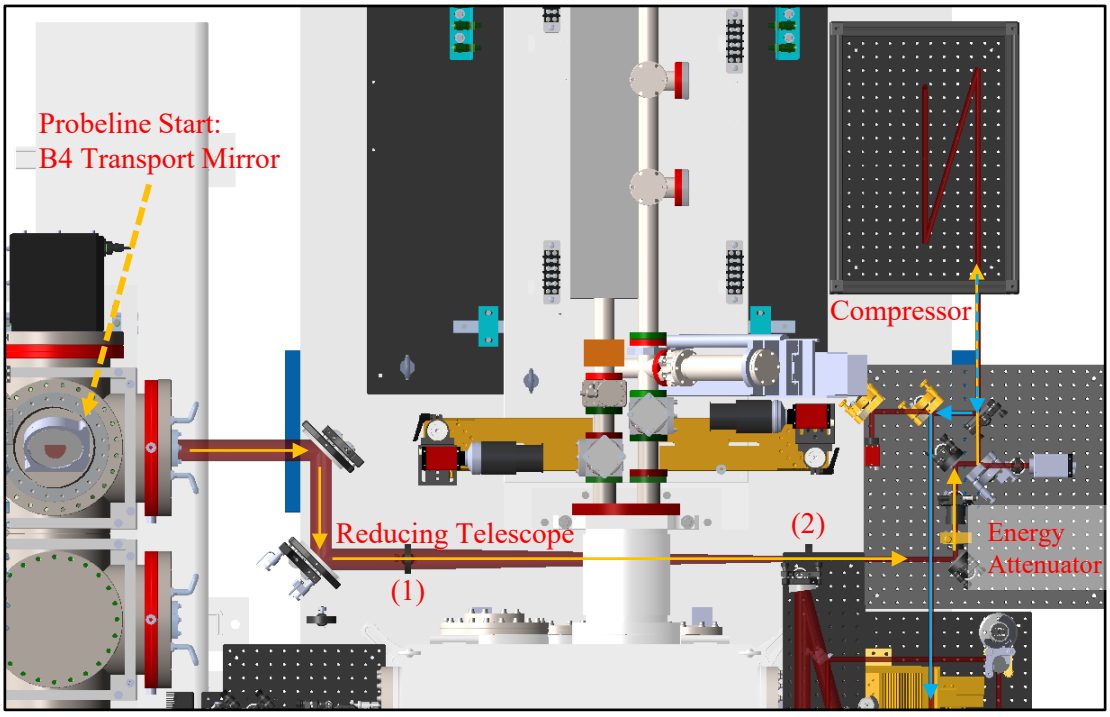

Figure 5.6: The initial section of probe laser line, orientation in this drawing is the electron beam travels from bottom to top. Starting with the leakage light transmitted through the B4 transport beam splitter (left side), the probe first passes through a telescope (lenses (1) and (2)) and reduces its size to 10 mm diameter (path shown in orange arrows). An energy attenuator is in place to remotely control the probe energy. The probe then passes into a compressor, which compresses the pulse length to < 70 fs. The probe then exits the compressor at a lower height and is sent upstream (path shown in blue arrows).

this section there are near-field alignment cameras before and after the reducing telescope, and a third after the compressor. Several mirrors are remotely motorized to allow for alignment while the electron beam is running.

The next section of the probe laser, depicted in the drawing of Fig. 5.7, represents the bulk of the external probeline optics. After the compressor, the beam is sent upstream (relative to the accelerator tunnel) to a delay stage that sets the relative timing of the entire probe to both the main laser line and the electron beam. The beam is then steered into a set of three beamsplitters that distribute the probeline energy across four separate probeline experiments and diagnostics. The primary beamsplitter reflects 80% of the laser energy downstream towards the “ionizer” and “shadowgraphy” probes while 20% is transmitted to the “EOS” and “shallow angle” probes. The

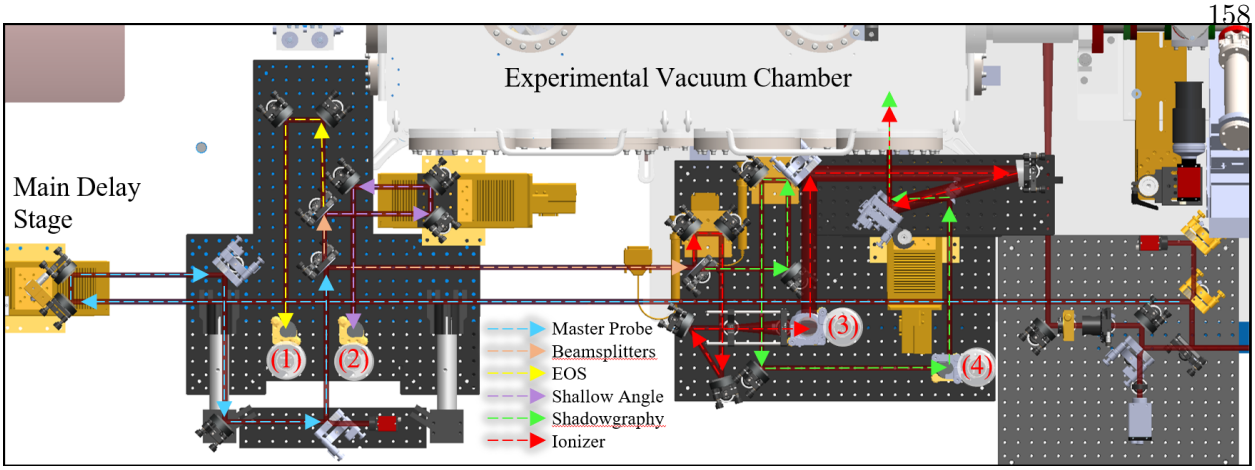

Figure 5.7: Drawing of the different probe lines in the experimental area. The master probe line continues from Fig. 5.6 on the bottom right in blue, passes through a main delay stage for the entire probe, and enters a series of beamsplitters. The line to the EOS is shown in yellow, the shallow angle in purple, the shadowgraphy in green, and the ionizer in red. At positions marked (1) - (4) periscopes send the laser to a higher height. The EOS and shallow angle probes continue on a breadboard above what is shown here before entering the experimental vacuum chamber.

secondary beamsplitters then further divide the laser energy. The reflected ionizer probeline receives 80% to the transmitted shadowgraphy probeline's 20%, while the beamsplitter for the EOS and shallow angle probelines is a $50-50\%$ beamsplitter.

While the ionizer probeline is the primary one of interest for plasma lens ionization, the other probeline arms can serve various experimental needs for a variety of experiments. As can be seen in Fig. 5.7, all probeline arms have their own delay stage to individually control their relative timing, with the exception of the EOS probeline arm which uses the master delay stage. The EOS-BPM allows for a non-destructive electron beam diagnostic for high precision measurements of the longitudinal charge distribution and transverse position. The shallow angle probe is a diagnostic for the longer plasma sources downstream of the experimental vacuum chamber, where the probe crosses the plasma at a shallow angle to pick up measurable phase distortions and learn about the density profile of the plasma. The shadowgraphy probeline is a similar concept, but for the short, high-density plasma formed in the gas jet outflow in the experimental vacuum chamber. This probeline arm was designed for the filamentation experiment[102, 103] which operates at a higher plasma density than the plasma lens experiment, so the plasma lens will not be using this

diagnostic for the time being.[6]

Now we turn our attention to the ionizer probeline. After the final beamsplitter the ionizer has access to an upwards of 20 mJ of laser energy. To prevent damage on optics we only send 10 mJ to the ionizer, which is sufficient for ionization given a good laser profile. A small delay stage allows for relative timing of the ionizer with respect to the overall probe timing from the main delay stage upstream. Two mirrors then steer the ionizer into an expanding telescope, which expands the ionizer to 30 mm diameter before sending the beam through the periscope labeled "(3)" in Fig. 5.7. The top mirror sends it towards the experimental vacuum chamber to optics on a higher optical stage shown in Fig. 5.8. Here, the top periscope mirror and the next mirror align the beam onto the center of the $f = 646$ mm OAP shown in Fig. 5.8. The beam reflects at a 15° angle onto a final steering mirror, which sends the beam into the experimental vacuum chamber.

Figure 5.9 shows the ionizer path through the experimental vacuum chamber. The laser is sent orthogonal to the electron beam's propagation axis and passes over the top of the gas jet. The gas jet is located on a 3D translation stage so it can be re-positioned as necessary. The ionizer's focus from the OAP is also aligned to the electron beam axis using the 150 mm translation stage in Fig. 5.8. This translation stage also has enough range of motion to scan the laser focus position and effectively change the ionized plasma lens thickness, as discussed in Sec. 4.1.3 and Fig. 4.6. There is additionally a 9.5 mm fused silica vacuum window flange between the final steering mirror and the gas jet target

After the ionizer passes through the experimental vacuum chamber, it exits on the far side through a vacuum window flange and enters an imaging system. Here, a beam sampler reflects 1% of the laser energy towards a camera while the remaining 99% is safely dumped. A single lens is installed to image the gas jet plane in a 2f imaging system. This both allows for an important diagnostic of the laser profile at and around the OAP focus, as well as one of the few sources of

[6] The filamentation experiment also employs a longitudinally-ionized plasma using the main laser line and an axilens focus, which we will use as a plasma lens in Sec. 5.2.5. However, since the shadowgraphy probe crosses at a transverse angle to the ionized plasma, the density needs to be quite high for a detectable phase advance in the shadowgraphy diagnostic.

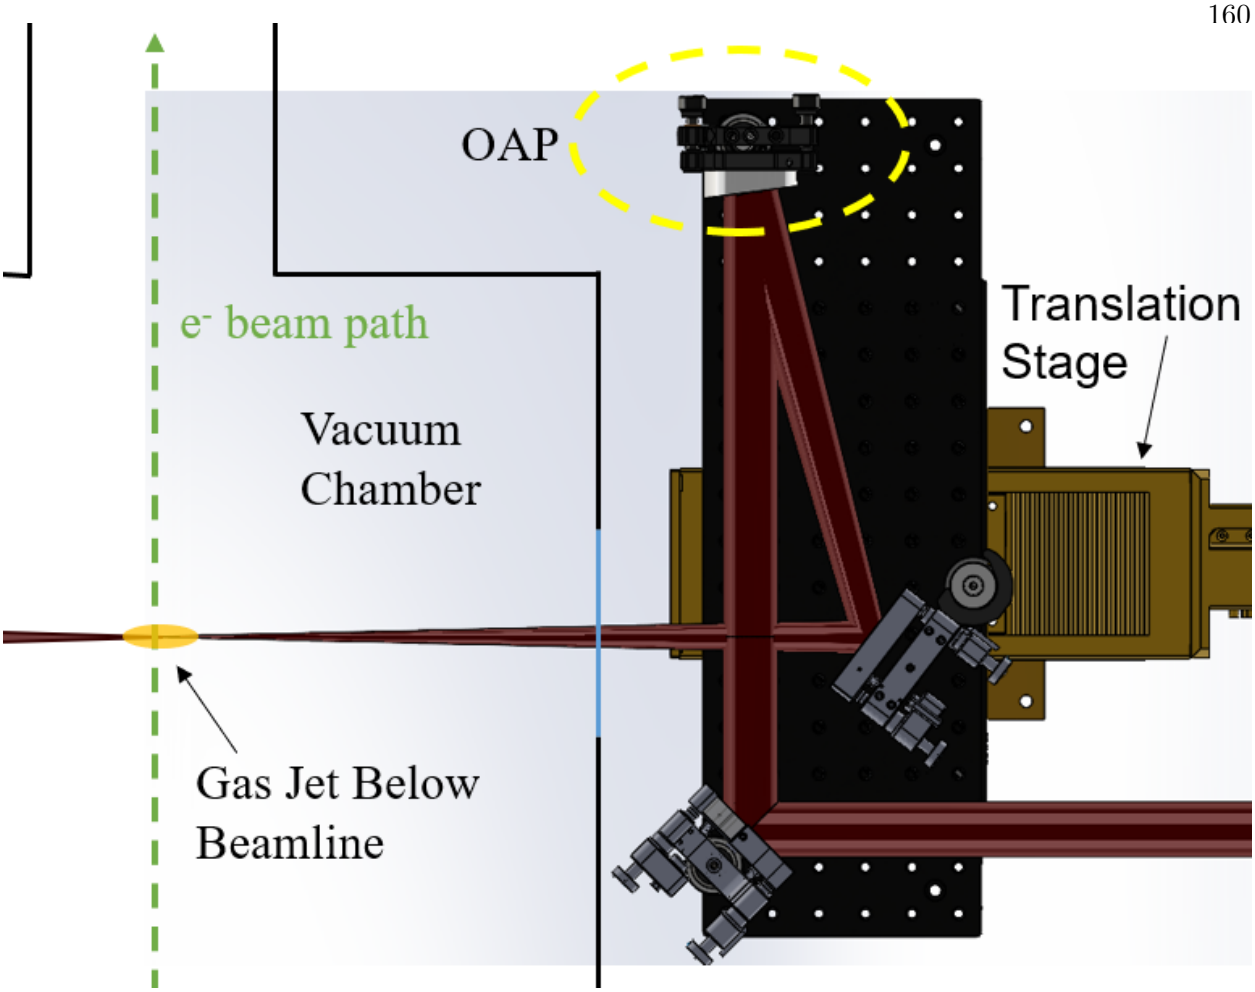

Figure 5.8: Drawing of the ionizer probe's stage immediately before the ionizer focus. The beam enters from the bottom-right after a periscope brings it to the breadboard height. Then a mirror allows for steering onto the center of the focusing OAP, which reflects the laser at a 15° angle. A final steering mirror then steers the beam onto the electron beam axis above the gas jet. The OAP and two neighboring mirrors are all placed on a 150 mm translation stage, which allows for the laser focus to be translated horizontally with respect to the electron beam axis. This is useful for aligning the focus position, as well as effectively scanning plasma lens thickness.

alignment feedback when attempting remote alignment of the ionizer probeline.

5.2.4 State of Transverse Plasma Lens Ionization

Plasma lens ionization using the transversely-propagating ionizer probe is the primary method we wish to pursue plasma lens experiments at FACET-II. Unfortunately, as of the time of writing we have been unable to generate plasma using this ionizer probe. There are a few issues which are

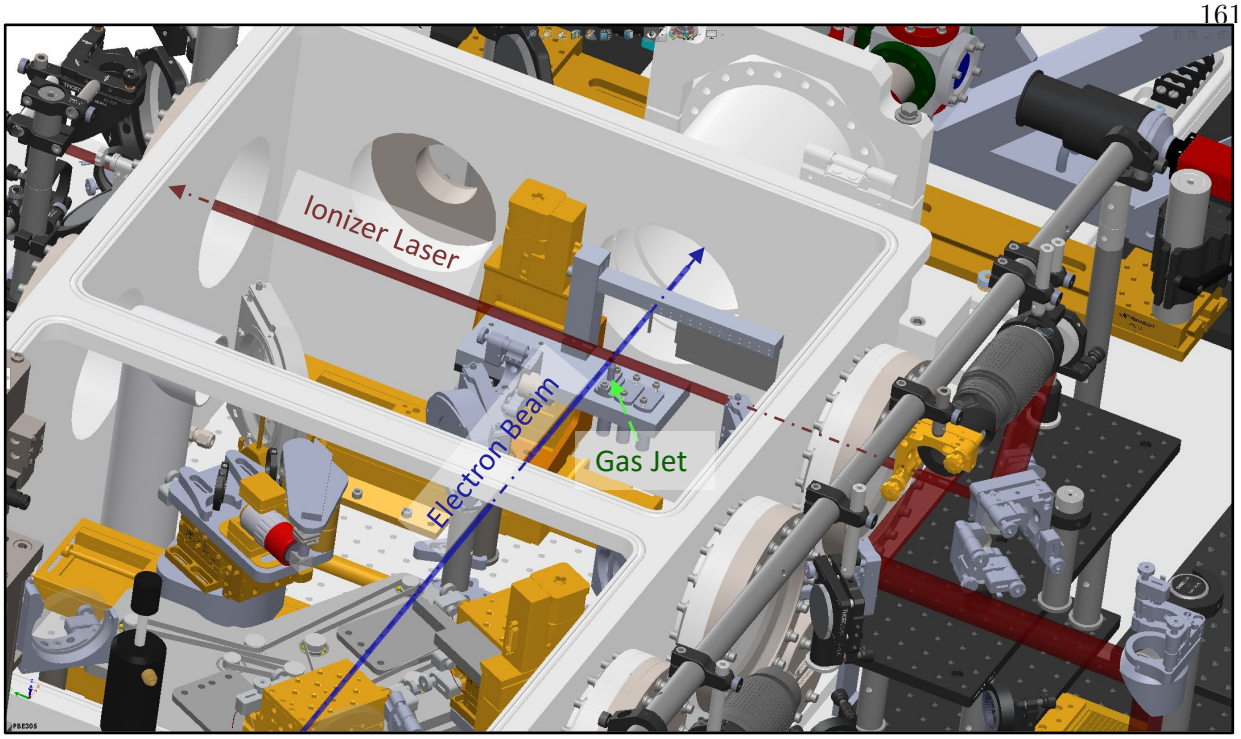

Figure 5.9: Drawing of the IP where the electron beam, ionizer probe laser, and gas jet outflow all intersect. The ionizer laser (red, right to left) is at a focus within the gas jet outflow (directed upwards) and ionizes a plasma lens, which then focuses the relativistic electron beam (blue, bottom to top). The gas jet is located on a translation stage with the capability for remote 3D motion.

currently preventing successful plasma ionization in the ionizer, all of which effect the quality of the focus at the gas jet outflow position.

The primary reason is simply that the wavefront quality of the laser at the ionizer probe is not good enough. As an example, Fig. 5.10 shows an image from the camera which images the ionizer's focus at the gas jet. Rather than a single, clean spot from the OAP focus, the focus shows signs of a poor wavefront with significant astigmatism. The quality of the laser wavefront for the FACET-II laser system is continually improving. For example, better procedures are being developed for more effectively using the DM to flatten the wavefront in the laser transport and main laser line. Such experimental procedures can eventually be extended to the complex probe line setup and will enable a clean wavefront for the ionizer.

The second issue is that of remote alignment. To achieve a good laser focus the beam must

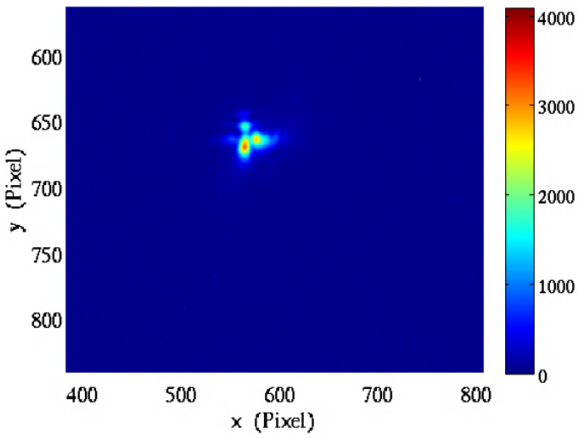

Figure 5.10: Camera image taken by the ionizer focus imaging camera. Here the camera is imaging the focus of the OAP on the plane of the gas jet. An ideal focus would be a single, clean laser intensity spot, but instead significant astigmatism and other wavefront issues leads to this multi-spot focus.

be carefully aligned through all of the optics so that it is not clipping on any mirrors and it is hitting the center of all lenses. This can be done during alignment shifts when personnel can be in the accelerator tunnel and manually align the full probeline. But once the tunnel is closed and the electron beam turns on, the significant temperature change causes optics to shift and alignment to be lost. There are only a handful of alignment cameras to aid in recovering good alignment. The entire probeline from B4 to the ionizer profile camera has twenty-four mirrors,[7] four lenses, the compressor, and an OAP; and the number of cameras for this setup is five.[8] The lack of alignment cameras before the OAP makes it especially difficult to remotely align the beam properly into the OAP without aberrations. In an attempt to make alignment easier we have tried replacing the OAP with a comparable transmissive lens that is more tolerant to misalignment, but this alone was unable to improve the focus quality enough to achieve ionization. An example of how poor alignment effects the laser intensity profile is shown in Fig. 5.11.[9]

The third issue is slightly more difficult to address through brute force alignment and laser work, and deals with temporal aberrations through the B-integral and nonlinear indices of refrac-

[7] Seven of which are motorized.

[8] Two for the initial reducing telescope, one after the OAP, one before the primary beamsplitter, and the ionizer profile camera at the end.

[9] The "C" stands for "Can you believe that this laser didn't ionize plasma?"

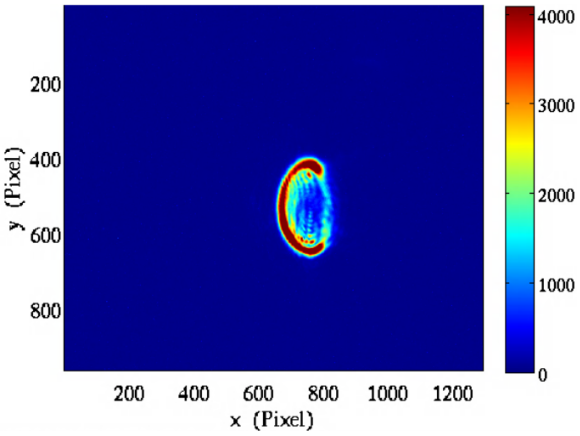

Figure 5.11: Camera image taken by the ionizer focus imaging camera. In this image, a plane upstream of the laser focus is shown. Significant clipping of the laser on various mirrors in the probeline result in this laser intensity pattern.

tion. When a high intensity laser pulse passes through a medium, there is a nonlinear phase shift which arises. This is quantified by the B-integral:

$$B = \frac{2\pi}{\lambda} \int n_2 I(z) dz \tag{5.1}$$

where λ is the laser wavelength, $I(z)$ is the peak laser intensity along the laser propagation axis, and n_2 is the nonlinear index of refraction for the material the laser is propagating through. Eqn. 5.1 represents the phase difference between the laser pulse center and edges, and a large B-integral indicates that self-focusing and self-phase modulation can occur. General practice is that you want to avoid a B-integral value above 1 to be safe from these aberrations [104]. For a laser propagating in vacuum, this nonlinear phase advance is not an issue except for any thick windows in the laser path. However, the FACET-II probeline is unfortunately almost entirely in atmospheric pressure air with a nonlinear index of refraction $n_2 = 3.01 \times 10^{-23} \mathrm{m}^2/\mathrm{W}$. For the full distance between the probe compressor and the expanding telescope for the ionizer, this 5.58 m of air propagation leads to a B-integral of $B = 4.81$. We also need to consider the transmissive optics of the ionizer's expanding telescope, and the 9.5 mm vacuum window flange immediately before the laser focus at the gas jet position. For the window flange, fused silica has a nonlinear index $n_2 = 2.19 \times 10^{-20} \mathrm{m}^2/\mathrm{W}$ and this leads to $B = 6.80$ using a 10 mJ pulse with 30 mm diameter. With all of these together, the

B-integral is well over unity.

For these reasons, we are currently unable to perform plasma lens experiments using the ionizer probe. This applies to the ionization schemes we investigated in Sec. 4.1.3 and Sec. 4.1.4, which both use an ionizing laser that propagates transversely to the electron beam axis. Fortunately, we are still able to perform some basic commissioning shifts by borrowing the experimental setup of the filamentation experiment, as will be discussed in Sec. 5.2.5. For the future, we are confident that the ionizer probe will eventually be useful. The first two issues can be solved with time as improvements are made to the full laser system and the probeline. The third issue with the B-integral is a little more uncertain, but it is possible to reconfigure the probeline to temporarily bypass the majority of the in-air propagation. This would allow us to test if a significant drop in the B-integral is sufficient to allow for ionization. Then, if needed, it could be possible to replace the vacuum window flange with a thinner alternative, make adjustments to the type of lenses used, cut the ionizer's laser energy to 5 or 1 mJ, or potentially do a partial redesign of the probeline layout.

5.2.5 Main Laser Line

While the ionizer probe is still being worked on, there is another method available for ionizing the gas jet outflow located in the experimental vacuum chamber. We can instead use the "main laser" and it's ~ 300 mJ of available energy. The main laser line continues down the vacuum laser transport after the B4 mirror that picks off 20% of the energy for the probeline arms. The remaining 80% is steered with final transport mirrors B5 and B6 into a large, in-vacuum compressor chamber (Fig. 5.12(a)). Here the laser is compressed to 55 fs before being sent towards the experimental vacuum chamber (Also known as the "Picnic Basket"). Alignment cameras after the compressor allow for this laser line to be remotely aligned through the compressor with good precision.

A gate valve with a window separates the compressor's vacuum chamber from that of the picnic basket (Fig. 5.12(b)). Here there are only three optics before the laser reaches the gas jet outflow at position (4). At position (1) is the axilens, which is mounted on a 3D translation stage

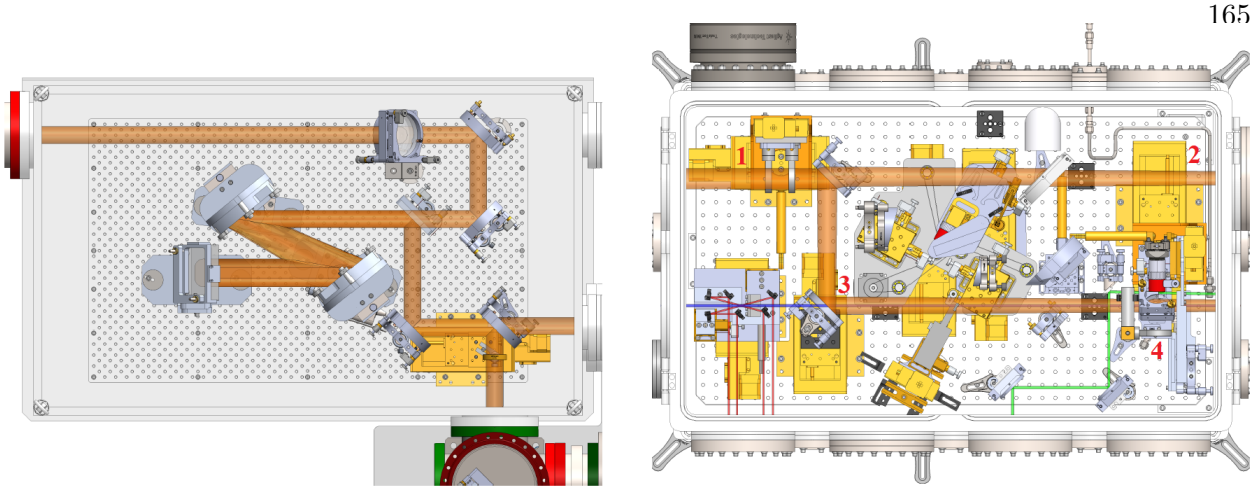

(a) Compressor chamber (b) Picnic basket

Figure 5.12: Drawing of the main laser line in the compressor chamber (a) and the main experimental vacuum chamber (b). The beam enters the compressor chamber in (a) from the laser transport tube in the top-left, and exits below in the bottom-right. A mirror sends the laser from the compressor chamber to the picnic basket in (b) at the top-left. Position (1) is the axilens mount, (2) is the pickoff for camera imaging, (3) is the upstream holed mirror, and (4) is the gas jet position. Reproduced from Ref. [33]

to allow for high precision alignment of the lens to the laser pulse. This axilens is a 2 inch optic with a 1 cm long line focus located at the gas jet position. A mirror afterwards then reflect the laser to a holed mirror at position (3), while a small fraction of the laser is transmitted towards a vacuum window flange at (2) to be imaged onto cameras for alignment and diagnostics.[10] The "upstream holed mirror" at (3) has a 5 mm diameter hole for electron beam clearance, and this mirror steers the laser so that it is co-propagating with the electron beam axis. At position (4) is the gas jet, and the focus of the axilens is also located at this position. Afterwards, the axilens focus divergences and is dumped on the walls of the electron beam tube further downstream.

Due to the efficiencies of the mirrors and optics in the laser transport and compressor, the laser energy that reaches the gas jet outflow is significantly less than the 800 mJ that comes out of the amplifier in the laser room. Only 609 mJ makes it to the B4 probeline beamsplitter, of this 487 mJ arrives at the main compressor. The main compressor has an efficiency of 72% and

[10] There are two alignment cameras, as well as a camera on a rail to image various positions along a long line focus. However, these cameras are set up for a different focusing optic with a focus further downstream and are incompatible with the axilens.

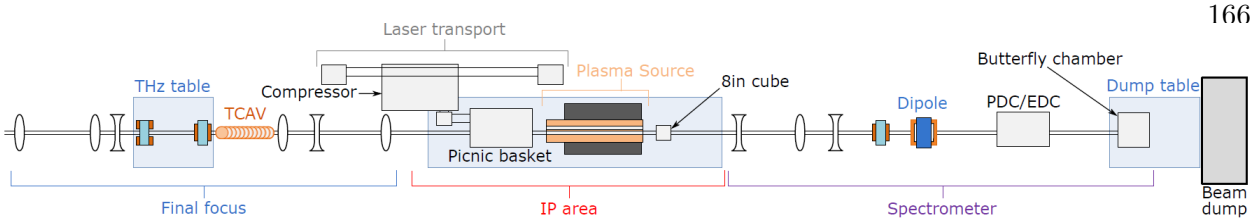

Figure 5.13: Top-view diagram of the layout of the FACET-II experimental area. From left to right is the final focus section with magnets for focusing the electron beam, the IP area section with the main experimental vacuum chamber (picnic basket) and plasma source chamber for PWFA experiments, and the spectrometer section with magnets for imaging the electron beam. The gas jet and plasma lens are located in the picnic basket. Reproduced from Ref. [33]

the laser energy after compression is 351 mJ. The final mirrors, optics, and windows between the compressor and the gas jet target further reduces this energy to a final value of 305 mJ. Still, this is orders of magnitude larger than what is available in the ionizer probe.

The higher laser energy for the main laser line makes it much more straightforward to ionize plasma than with the ionizer probe. There are also significantly less mirrors to align through, plenty alignment cameras to assist with remote alignment, and the entire main laser is under vacuum. The main downside arises from the fact that the laser co-propagates with the electron beam axis. The axilens focus is a line focus a few mm's long, and so the entire gas jet outflow is ionized along the laser axis. This results in a plasma lens with a thickness on the order of the gas jet nozzle exit. So, while the ionizer probe can transversely ionize a sharp plasma density profile with thicknesses on the order of 100′s μm, the main laser ionizes a plasma lens with a smooth density profile and thicknesses many mm's long.

5.3 Experimental Area

In this section we take a look at other components in the experimental area apart from the ionizing lasers. This includes the gas jet which was installed at the time of plasma lens commissioning, a camera which images the ionized plasma from above, and the downstream electron beam spectrometer. A diagram of the experimental area is shown in Fig. 5.13.

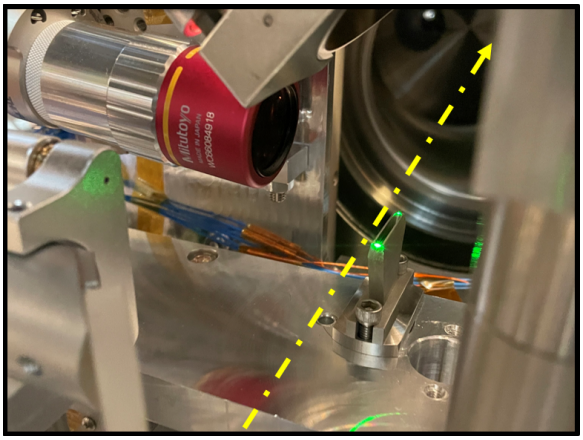

Figure 5.14: Picture of the elongated gas jet in the experimental vacuum chamber. The electron beam axis is shown in dashed yellow, and this is also the main laser path for the axilens focus. The microscope objective is used with the shadowgraphy probe.

5.3.1 Gas Jet and Top View

As motivated earlier in Chapter 4, a gas jet outflow is an ideal location to generate a plasma lens. The limited volume of the neutral gas makes it easier to fully ionize the plasma lens without effects of laser refraction, and having less gas in the accelerator is generally better for the electron beam and accelerator equipment. The gas jet at FACET-II is located within the picnic basket on a target mount stage with 3D translational motion. This target mount stage is designed to accommodate multiple experiments, with the ability to host several solid targets as well as up to three gas jet nozzles. As of time of writing, only one of the three nozzles is installed[11] but the exact nozzle can be swapped during tunnel accesses. The gas jet is supplied by high pressure gas line from the Klystron gallery that is capable of supplying up to 1000 psi of Hydrogen or Argon gas, and a Parker Iota One controls the opening of a solenoid pulse valve to be a few ms long.

Figure 5.14 is a picture of the gas jet nozzle as it is installed in the picnic basket. The available nozzle at the time of plasma lens commissioning was an elongated nozzle used for the filamentation experiment and a much higher density gas outflow. This nozzle opening is 2 cm long in the direction of the electron beam axis. For a plasma lens experiment we compensate for the

[11] Getting a high-pressure gas line to be leak-proof in a vacuum chamber is difficult enough for a single gas line, and this gets exponentially more tricky when adding a manifold with three gas jets in close proximity.

long gas jet nozzle by decreasing the backing pressure to lower the density of the gas in the outflow.

To be able to observe the ionized plasma, a camera is installed which images the ionized plasma from above. In Fig. 5.14 a mirror is partially visible above the gas jet. This mirror allows for a camera outside of the picnic basket to view the plasma. This external camera can be seen in Fig. 5.9 above the ionizer's OAP.

The neutral gas supplied by the gas jet is pumped out through a turbo pump connected to the experimental vacuum chamber. This pump is sufficient to adequately drain enough gas between electron beam shots to operate the plasma lens at 10 Hz for relevant backing pressures. A differential pumping system also separates the experimental vacuum chamber from accelerator sections further upstream and downstream. This system uses a series of narrow apertures and vacuum pumps to gradually transition from vacuum conditions to non-vacuum conditions. These narrow apertures allow for the electron beam to travel from vacuum conditions upstream into the experimental vacuum chamber without passing through any solid windows.

5.3.2 OTR Screens and Alignment Laser

With the electron beam propagating through the experimental area, it is useful to know the exact trajectory of the electron beam axis so that the experimental apparatuses and high power lasers can be aligned to this axis. To achieve this, there are a number of retractable metal foils along the electron beamline in the IP area of Fig. 5.13. These foils emit optical transition radiation (OTR) light as the electron beam pass through, and each OTR foil has a camera which images the surface of the foil. OTR light is emitted due to complex field dynamics within a metal as an electron beam approaches it [105]. The result is a highly-directional yield of photons that scales with the number of electrons and not with the energy of the electrons. Using these OTR foils we can grab reference positions on the cameras for where the electron beam axis is.

Just before the main experimental vacuum chamber, a mirror can be inserted that directs a low-energy, green, HeNe laser[12] into the vacuum chamber along the electron beam axis through

[12] Affectionately known as "GreeNe"

the entire IP area. This laser can then be steered onto the same reference positions as the electron beam to get an accurate and visible representation of the electron beam axis. For example, this is useful when aligning the gas jet. In Fig. 5.14, the green mark on the front edge of the gas jet is from this green alignment laser. This lets us know the relative positions of the gas jet and electron beamline to a high degree.

5.3.3 Electron Beam Spectrometer

The primary tool by which to measure the electron beam's interaction with the ionized plasma is to image the electron beam onto downstream screens using the FACET-II electron beam spectrometer. This spectrometer is made up of three quadrupole magnets that form an imaging system for the electron beam in the horizontal plane, and a dipole magnet that disperses the electron beam vertically with respect to electron energy. To maximize the energy resolution, the vertical size of the electron beam is focused to a small beta at the image plane. The spectrometer can be tuned to allow for the object plane's longitudinal position to be varied, which allows for scans of the electron beam's size around the gas jet position.

There are two types of screens we can use to view the electron beam: a scintillating YAG screen for low charge-density beams and a polished titanium screen to view OTR light emitted as the electron beam passes through. For our purposes here, we will use FACET-II's standard 500 μm thick polished titanium for OTR measurements. Two cameras view this screen: DTOTR1 has a small field of view with a narrow range of 2 GeV visible electron spectra, while DTOTR2 has a large field of view and can observe the electron beam's full spectra from $0 - 10$ GeV. From the diagram on Fig. 5.13, this screen is located at the butterfly chamber near the right side.

In addition to only being able to image the horizontal beam size, the spectrometer can also only image a single energy slice at a time onto the OTR screen. Energy slices outside of the imaged energy slice will appear larger on the screen. The imaging of the electron beam using this electron beam spectrometer will be discussed in greater detail in Chapter 6 when it is applied towards the plasma lens commissioning shift.

Chapter 6

Commissioning Data

While there are many positive aspects with working on an experiment at a large facility like FACET-II, there are additionally a few downsides. These include delays in getting the facility to full operation,[1] small windows of time to perform installation and alignment work, and sharing experimental priorities with the many of other experiments at FACET-II. But this also one of FACET-II's greatest strengths; with many collaborators working together there are opportunities to improvise and assist with the experimental needs of each other. For the case of the plasma lens, when the transverse ionization using the probeline was not working as intended, we were fortunate enough that another experiment set up in parallel could be slightly modified to be used for early plasma lens commissioning. For this, we are grateful to the E305 filamentation experiment for allowing us to borrow their axilens and elongated gas jet setup.

In this chapter we present the findings from the first plasma lens commissioning shift. With access to only a single electron bunch, we are operating in the single bunch regime of Sec. 3.6.3. This electron beam is also much lower charge density than FACET-II's expected beam density once the machine is fully tuned. With this lower beam density comes a smaller parameter space where the plasma is operating in the nonlinear blowout regime. Rather than testing proper thin plasma lens operation, this shift is more about colliding an electron beam with a plasma to try and observe focusing of the beam. In the process, we gain experience with the shift procedure and data analysis tools. It is far from ideal in the context of what will be needed in the future, but it is a good first

[1] Global pandemics don't help, either.

step in testing the setup and learning what the experiment needs going forward.

6.1 Data Collection

We begin by describing the procedure of the commissioning shift. We first set up the experimental area by aligning the laser and gas jet to the electron beam axis. The accelerator operators during this time are tuning up the beam parameters to the desired parameters as close as possible. Once everything is set up, data acquisition during the shift is described in Sec. 6.1.2.

6.1.1 Shift Setup

For the plasma lens experiment, we need to have the electron beam, gas jet outflow, and ionizing laser all aligned to each other spatially and temporally. And since both the gas jet and laser can be remotely aligned in the sector 20 experimental area, the first step is then to get a reference orbit for the electron beam. First, the operations team at FACET-II tune up the electron beam to parameters as close to the desired parameters as possible (Fig. 6.1). As this shift was taken in the commissioning phase of FACET-II, the beam parameters are less than ideal but sufficient enough for the early goals of the plasma lens experiment.

The accelerator operators measure the electron beam parameters before and after the shift, and we use these measurements to get an idea for the average parameters during the shift. The beam charge was approximately constant at 1.5 nC. The beam energy slowly drifted from 10.10 to 9.87 GeV during the shift, so we assume an approximate energy of 10 GeV. The transverse beam sizes are measured using a wirescanner[2] measurement before and after the time of the shift,[3] the longitudinal beam size is measured using a transverse deflecting cavity (TCAV) measurement,[4] and

[2] Wirescanners are a 1D measurement of the beam size over many shots. The operators move a thin wire across the electron beam's axis and measure the amount scattered electrons. From this, one can infer the transverse beam size in the scanned direction.

[3] As the wirescanners are located slightly upstream of the experimental vacuum chamber, the standard practice is to use the final focus quadrupoles to slightly shift the focus upstream onto the wirescanner position. This allows for the focus to be measured and the spot size difference between the shifted and unshifted settings is minimal.

[4] TCAVs are a destructive measurement of the beam's longitudinal size. It works by using an X-band wave within the TCAV that transversely deflects electrons based on their longitudinal position. A downstream monitor can then correlate the deflection angles with the longitudinal beam size.

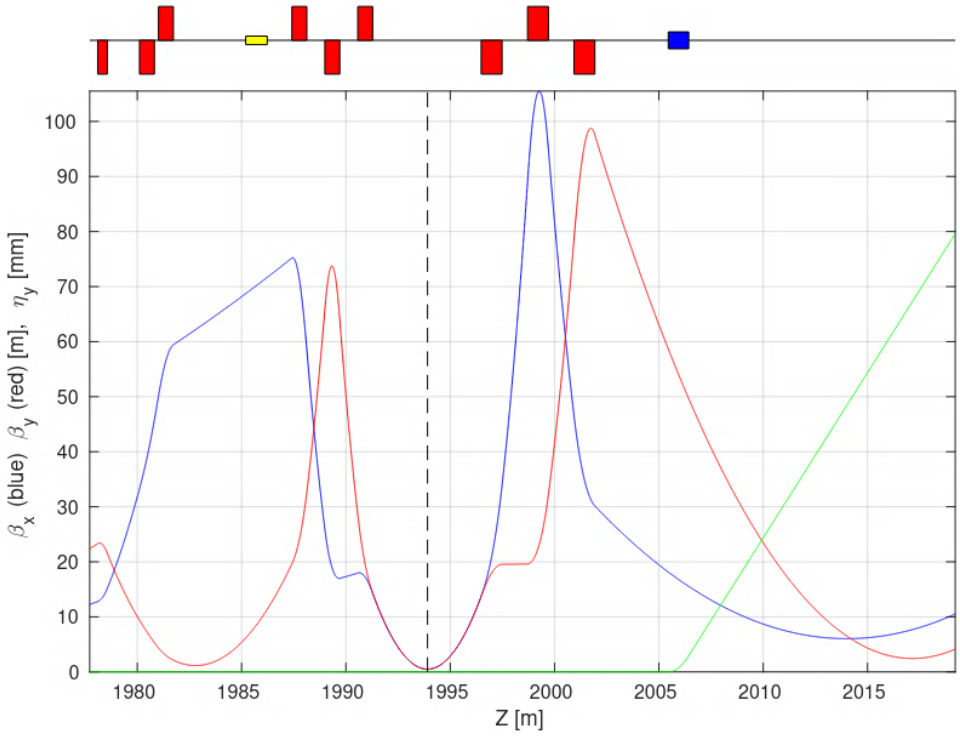

Figure 6.1: Target configuration for the electron beam betafunctions (blue for β_x, red for β_y) as determined by the Sector 20 quadrupole (red rectangles, above) settings. The focus at the gas jet plane (dashed) is designed for a 50×50 cm betafunction focus.

the emittance is measurement by taking a multi-wirescanner measurement at the start of L3 [106]. These values are averaged to determine approximate beam parameters during the shift, and are summarized in Table 6.1.

Once the electron beam orbit through the accelerator is locked in, two OTR screens in the experimental area are moved into the beamline: one at the position of the gas jet on the target mount and another approximately 50 cm downstream. Reference positions of the OTR light are then taken by cameras that image the surface of these foils.

After the reference positions are taken, we block the electron beam from entering the experimental area and align the low energy green alignment HeNe to the reference marks on the experimental area OTR screens. This sets the green HeNe on the same axis as the electron beam. We then move the gas jet so that we can see the green HeNe reflecting off of the center of the top

Table 6.1: Electron beam parameters during the plasma lens commissioning shift, as determined through linac diagnostics taken by the accelerator operators before and after the shift. The peak density and betafunctions are calculated assuming a Gaussian beam distribution.

Property	**Electron Beam**
Charge	1.5 nC
Energy	10 GeV
ϵ_{nx}	23.63 ± 18.23 μm rad
ϵ_{ny}	10.46 ± 2.86 μm rad
σ_x	38.55 ± 1.69 μm
σ_y	32.50 ± 2.25 μm
σ_z	27.86 ± 5.12 μm
$\sim n_{b,peak}$	$1.70 \pm 0.343 \times 10^{16}$ cm^{-3}
$\sim \beta_{x,calc}$	1.23 ± 0.956 m
$\sim \beta_{y,calc}$	1.98 ± 0.606 m
$\beta_{(x,y),target}$	0.5×0.5 m

of the nozzle (As depicted in Fig. 5.14). This lets us mark the motor positions of where the gas jet is aligned to the electron beam, and allows us to move the gas jet 3 mm downwards when we send in the electron beam for experiments.

Similarly, we can use the electron beam reference marks to align the longitudinally propagating main laser line. First, we align the laser through the Sector 20 laser room, laser transport, and main compressor chamber as described in Chapter 5. Once the laser is in the main experimental vacuum chamber, we use the two mirrors upstream of the gas jet to align the laser onto the marks on the OTR screens. After this is done, we move the axilens into the main laser line in the experimental vacuum chamber. Fine-tuning of the axilens alignment is done by moving the lens until the focus is also observable on the reference mark on the OTR foil near the gas jet location. With the axilens aligned to the electron beam axis, we can now optimize the ionized plasma by pulsing the gas jet and observing the plasma glow imaged by the TopView camera. Here we primarily want to shift the longitudinal position of the axilens such that the ionized plasma is centered longitudinally in the elongated gas jet. After this alignment, images of the plasma glow from TopView look like those in Fig. 6.2.

Temporally, the laser and gas jet are timed into electron beam arrival time by using the signal

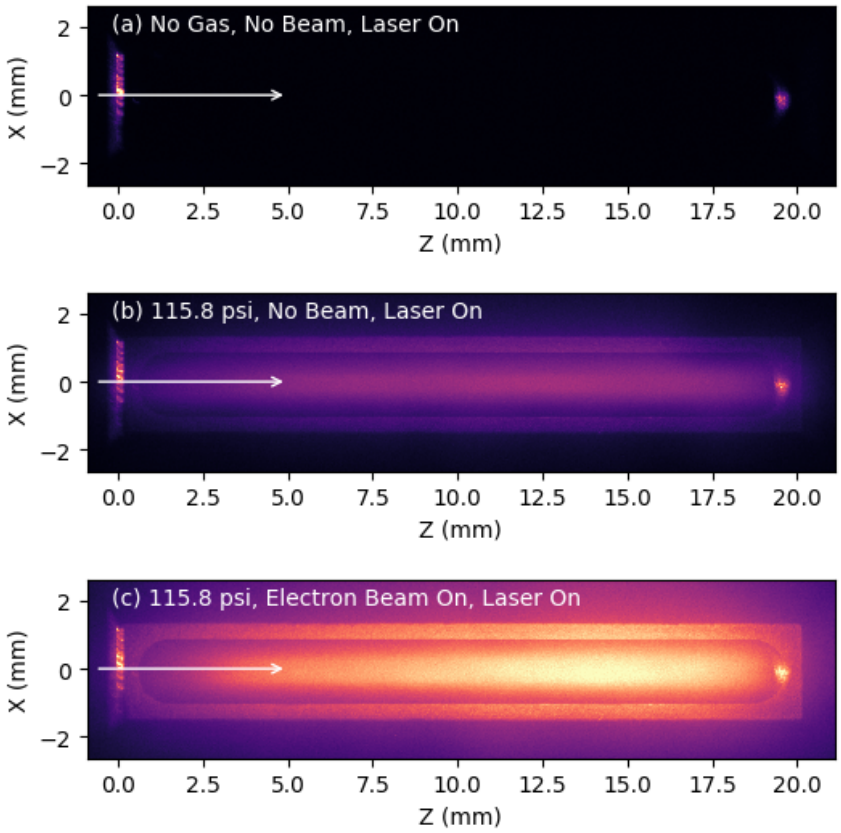

Figure 6.2: Images taken by the TopView camera during the shift setup phase, laser and electron beam axis is shown as the white arrow. (a) is an image taken when the laser is passing above the gas jet but there is no gas in the chamber. Reflected laser light off of the gas jet nozzle is observed. (b) is an image with the gas jet pulsing with at backing pressure 115.8 psi, resulting in a visible plasma. (c) is an image with additionally the electron beam passing through. The electron beam deposits significant energy into the plasma and the light emitted is increased.

from the plasma itself. While the timing has mostly been found in previous general commissioning shifts, we can fine tune the alignment here by shifting the laser's arrival time with respect to the electron beam's master RF signal. If we are watching a downstream diagnostic for the electron beam, there will be a clearly visible distortion in the electron beam if the laser arrives before the electron beam and ionizes a plasma. If the laser arrives after the electron beam, the downstream diagnostic will show an unperturbed beam. By repeating this search iteratively with smaller temporal windows, it is relatively easy to time the laser to the electron beam with sub-ps precision. This is sufficient because on the lower end the laser arrival time jitters on the order of 10 fs while

on the upper end the plasma expands and decays on the timescale of a few ns. For the gas jet, since the time scale of gas motion is much longer (order of ms) the tolerance for temporal alignment is very large. It can still be tuned up by watching the TopView camera and maximizing the brightness of the ionized camera when shifting the gas jet valve's opening signal.

During the shift the gas jet settings were set to be a 10 Hz repetition rate (same as the electron beam's repetition rate) with the opening time on the valve set to 2 ms. The gas jet is positioned so that the nozzle exit is 3 mm below the electron beamline. We use diatomic Hydrogen gas, as the lower ionization energy would make it easier to ionize a plasma. A manual pressure regulator allows us to control the backing pressure for the gas jet, and this regulator has a range that lets us scan from 1 to 116 psi. This particular elongated gas jet has previously been characterized to have a density profile with $n = 2.7 \times 10^{18}$ cm^{-3} with a backing pressure of 1000 psi[5] , so assuming a rough linear scaling this would imply our backing pressure range corresponds to a gas density range of 2.7×10^{15} cm^{-3} to 3.15×10^{17} cm^{-3}.

6.1.2 Shift Procedure

The bulk of this commissioning shift involves using the electron beam spectrometer to image the electron beam's horizontal size along with its energy spectrum. We take datasets at various backing pressures to vary the plasma density while keeping the laser and electron beam parameters constant. For each dataset, we scan the spectrometer's object plane in a region of 2 m around the position of the gas jet. This region is split into 20 steps, and at each step 10 shots are recorded. The one exception is the dataset taken with zero backing pressure, in which only 7 steps are taken. These datasets are summarized in Table 6.2. Due to the limited range of available backing pressures, comparing with the beam parameters of Table 6.1 reveals that electron beam ($n_{b,peak} \approx 1.7 \times 10^{16}$ cm^{-3}) is only more dense than the gas jet outflow up to a backing pressure of $\sim$ 6 psi. This comparison assumes a lot about the beam and gas jet,[6] but we can expect

[5] From private correspondence with Choajie Zhang of UCLA and Alexander Knetsch of Ecole Polychnique

[6] First, that the peak density of the gas jet outflow linear scales with backing pressure across three orders of magnitude. Second, while the electron beam's transverse profiles are approximately Gaussian, the longitudinal

Table 6.2: Summary of the main datasets which were taken during the commissioning shift and used in the analysis of this Chapter. The peak plasma density and corresponding plasma wavelength is estimated from previous characterizations of the same gas jet. The step size refers to the object plane scan from 1.27 m longitudinally before the gas jet to 0.73 m after the gas jet.

Backing Pressure	0 psi	1 psi	6 psi	24 psi	57.8 psi	115.8 psi
Approx. Den. ($\times 10^{16}\,\mathrm{cm}^{-3}$)	0	0.27	1.62	6.48	15.3	31.5
Plasma Wavelength λ_p (μm)		102	41.7	20.9	13.6	9.46
Longitudinal Step Size (cm)	33	10	10	10	10	10

that the datasets with a higher backing pressure will almost certainly be operating in the linear regime rather than the nonlinear blowout regime. We also note that the transverse beam size from Table 6.1 becomes comparable to the plasma wavelength with a backing pressure of 6 psi. The plasma wavelength is the approximate length scale of blowout plasma wakes, and so even at 1 psi it could be possible that the with a statistical rms beam size of 38.55 μm the full transverse width of the beam can be larger than the wake sheath's blowout radius.

During the datasets, magnet settings for the spectrometer are varied according to a script which calculates necessary magnet strengths to image the horizontal electron beam size at the object plane onto the image plane at the downstream OTR foil (DTOTR2). This is calculated by using the transfer matrix of the electron beam from the gas jet position, through the spectrometer magnets, and onto the DTOTR2 position. This transfer matrix can be calculated using the magnet positions and thicknesses on Table 6.3 with the general transfer matrix for a thick quadrupole with drift spaces in between (See Eqn. 2.12 in Sec. 2.1.2). We can also use this transfer matrix to calculate the beam magnification at DTOTR2. As the magnet settings change during the dataset, the magnification also changes during the dataset. For a position 1.27 m upstream of the gas jet the object plane magnification is 2.43, while 0.73 m downstream of the gas jet the magnification is 3.61.

For each dataset we record at every step the electron beam spectrum on DTOTR2, the plasma glow as seen on the TopView imaging camera, and a collection of scalar values from toroids

profile can exhibit density spikes and other irregular features.

Table 6.3: Summary of the positions of the three magnets (Q0, Q1, Q2) in the imaging spectrometer and the downstream OTR position with respect to the gas jet location. Magnet positions refer to the center of the magnet, and all three magnets are 1 m in length.

	Gas Jet	Q0	Q1	Q2	DTOTR2
Longitudinal Position	0 m	3.71 m	5.94 m	8.16 m	21.99 m
Thickness		1 m	1 m	1 m	

and beam position monitors in the experimental area that record the electron beam charge and transverse position. The resolution of the DTOTR2 camera is 30.3 μm/px. DTOTR2 is originally meant to be a low resolution, high field-of-view image of the spectrum, while its high resolution counterpart, DTOTR1, was unusable during this shift. The downstream OTR foil is oriented so that its face is at an angle to the electron beam axis. DTOTR1 looks at the reflected light off of this surface while DTOTR2 looks at the transmitted light on the back side. Unfortunately, during this shift the electron beam was bright enough to saturate the DTOTR1 camera and so we are left with only DTOTR2.[7]

For each dataset, the data acquisition software at FACET-II initially records a single shot with the electron beam blocked to act as a background for all of the diagnostics. During this background shot, the laser is still on and so we can use this background shot for the TopView camera to get an image of the ionized plasma glow without a perturbation from the electron beam.

6.2 Data Analysis

Now we summarize the analysis we have performed on data taken from the commissioning shift. First we look at the TopView camera images to find the approximate shape of the plasma density profile. Next we take the DTOTR2 beam spectrum images to find the evolution of the horizontal beam size as the object plan is varied. We fit this evolution to that of a beam propagating in free space, and we perform this fit for both the beam's full projection as well as with separate energy slices of the beam. Lastly, using toroid measurements and intensity on DTOTR2, we investigate if there are any signs of charge loss in the electron beam.

[7] This has since been corrected. One of many diagnostics unprepared for the sheer brightness of these beams.

6.2.1 Top View of Ionized Plasma Glow

First we take a look at the TopView camera images which image the light of the ionized plasma from above the gas jet. This can give us information on roughly what the plasma density profile is for the ionized plasma. The TopView camera has a resolution of 17.94 μm/pix and an exposure time of 100 μs. It is important to keep in mind that this by itself is not a true density diagnostic. The plasma glow recorded by the camera is an aggregate of the recombination process and collisional excitation of neutrals over a large window of time, with photons that reach the camera having originated from generally unknown positions in the plasma. Still, since a higher density generally yields more photons, this diagnostic is still useful for a relative comparison between the datasets.

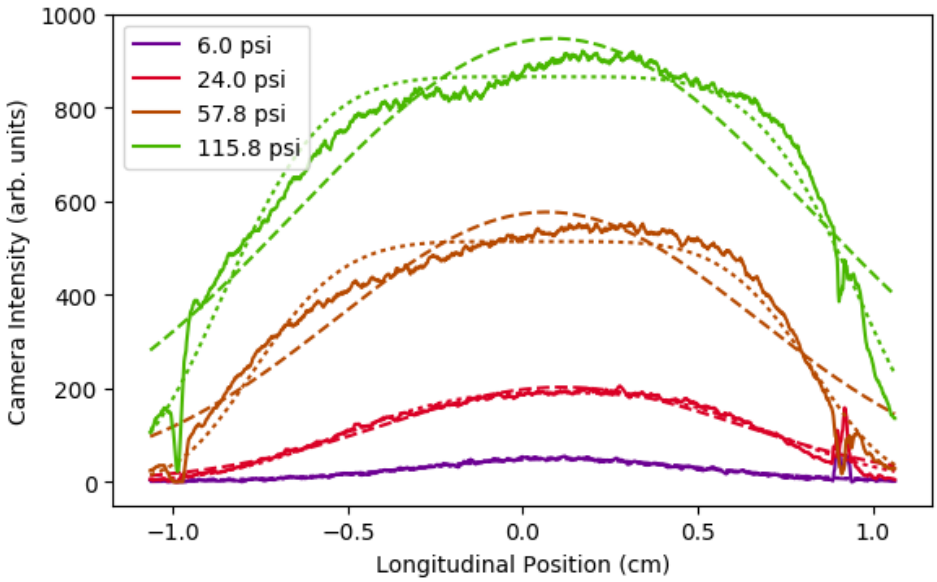

Figure 6.3: Lineouts of the plasma glow visible on the TopView camera diagnostics. Lineouts are taken along the longitudinal center of the gas jet outflow on the electron beam axis and background-subtracted for the case with 0 psi. The 1 psi case is not plotted here because the signal was too low and noise-dominated. Each curve shows the plasma glow intensity with an ionizing laser but not a co-propagating electron beam. The spiky features on either end are the edges of the gas jet exit that become illuminated from plasma glow, and remain even after background subtraction. Dashed lines are a standard Gaussian fit that works well for the lower density cases, dotted line is a Super-Gaussian fit that works better at the higher densities.

We analyze a longitudinal slice along the electron beam axis for the background shots where the laser ionized the plasma and the electron beam was not sent through (Fig. 6.3). Here we get a

general idea of how the plasma density was distributed along the length of the elongated gas jet. The laser and plasma glow illuminate the edges of the gas jet exit, so the gas jet outline appears in this figure despite background subtraction. At 1 psi the signal was weak and noise-dominated, so the profile is assumed to be similar to that of 6 psi. The shape of the plasma glow does qualitatively change as the backing pressure is increased. At the low backing pressures, the gas outflow is more centralized in the middle of the elongated nozzle. This makes sense, as the nozzle is designed for experiments operating at a much higher backing pressure and so a higher backing pressure may be needed for the gas to evenly populate the outflow region.

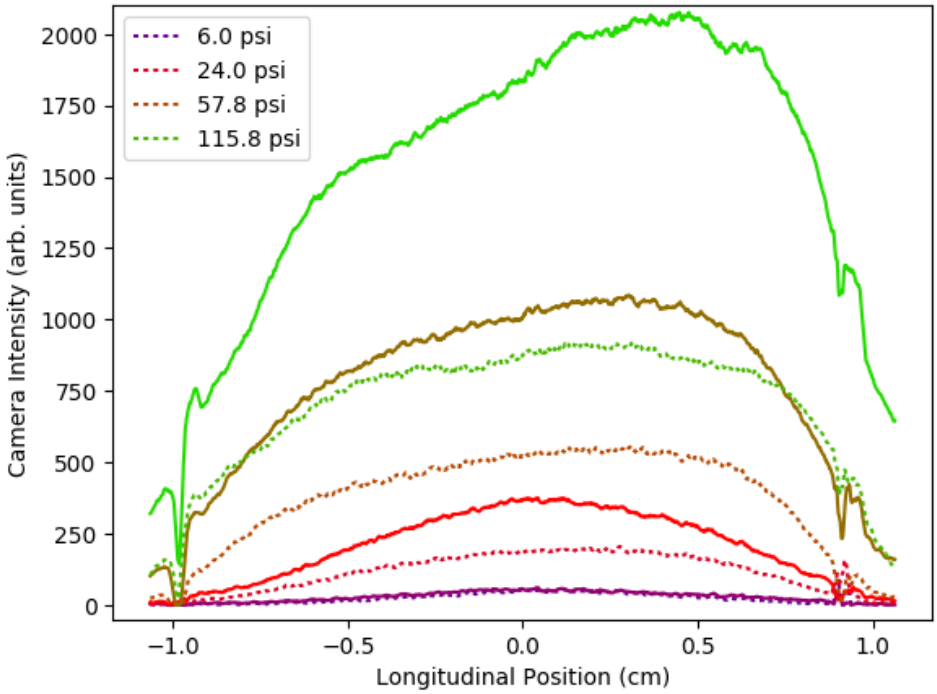

Figure 6.4: Lineouts of the plasma glow visible on the TopView camera diagnostics. The dotted lines are for no electron beam (Same as Fig 6.3), while the solid line is the plasma glow when the electron beam is additionally sent through the plasma.

Next, we compare the these longitudinal plasma glow slices for the cases with and without an electron beam. As mentioned before, the data in Fig. 6.3 is taken as background images to their respective datasets that include the electron beam. Figure 6.4 shows longitudinal lineouts of the plasma glow with and without an electron beam. There is a significant increase in the amount of

light collected by the camera when an electron beam is propagating through.

To gain a better idea of if this increase in plasma glow signal is due to beam ionization, we have two additional datasets that were taken in the setup phase of this commissioning shift. These datasets include the pulsing gas jet and the electron beam, but do not include an ionizing laser or an object plane scan for the spectrometer. The intensity lineouts from TopView for these datasets are shown in Fig. 6.5, and do show a signal for beam ionization at 150 psi backing pressure. However, this plot uses the same vertical scale as before, so the signal here from beam ionization alone is about $< 5\%$ that of laser ionization when compared with the 115.8 psi case of Fig. 6.3.[8] While beam ionization could be a small factor in increasing plasma density, the significant boost in signal of Fig. 6.4 is more likely due to the electron beam depositing energy and inducing more collisions in the pre-ionized plasma.

The negligible effect of beam ionization on the plasma density profile is a nice result for both plasma lens design and for modelling the plasma from this commissioning shift. The plasma density profile is primarily determined by the ionizing laser's intensity profile, and so it is possible to design a plasma lens with a more defined volume without beam ionization significantly altering the density. While this commissioning shift has a relatively wide electron beam with low density, one can also choose gas species with a higher ionization energy if beam ionization is a problem. For the purposes of analyzing the density profile of this shift, we can just assume the density profile is that of when only the ionizing laser is propagating through the gas jet.

If we take the camera intensity lineouts for the data with laser ionization and no electron beam propagation, we can take super-Gaussian fits to the profile and estimate the effective plasma length by integrating the fit. These fits are shown in Fig. 6.3, and the effective plasma lens dimensions are summarized in Table 6.4 assuming an ideal underdense plasma lens. Even if the nonlinear, underdense regime was valid for all densities, the normalized thickness $\sqrt{K}L$ is larger than 1 for all but the lowest density.[9] Only at 1 psi can our plasma source be labeled as a "thick lens."

[8] 5% signal doesn't necessarily mean a 5% ionization rate due to the beam. The beam could be fully ionizing in a very small region, which results in a total light on the camera which is 5% that of the laser's ionization signal.

[9] Since the transfer matrices for a thick lens depend on $\sin\sqrt{K}L$, for large values of $\sqrt{K}L$ an electron beam

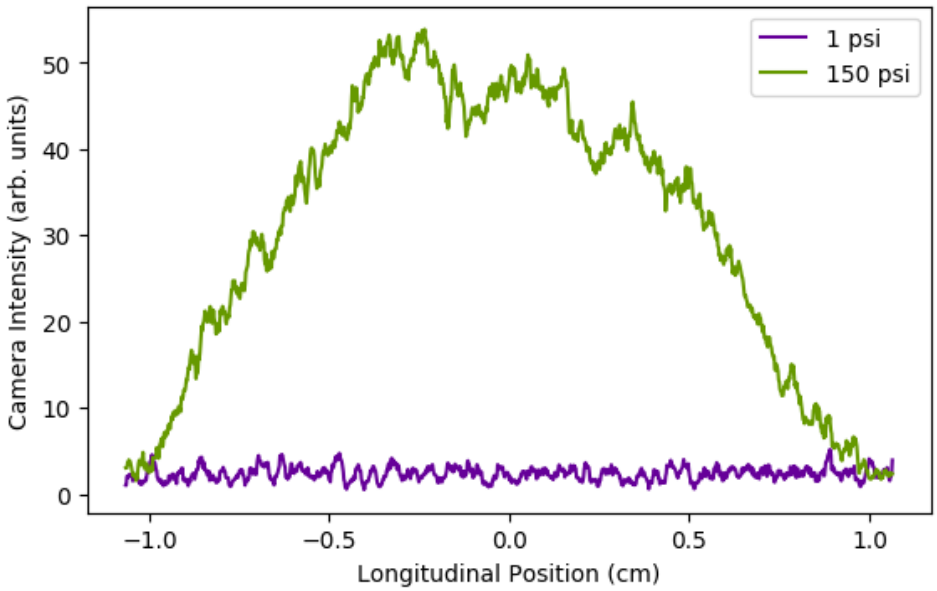

Figure 6.5: Similar to Fig. 6.3, lineouts of the plasma glow visible on the TopView camera diagnostics. Plotted here are the intensity lineouts for two datasets without an ionizing laser using the same vertical scale as before. At 150 psi is a small, but detectable ionized plasma due to the electron beam fields. Comparing the vertical scales from before, this ionization signal is about $< 5\%$ that of the signal from laser ionization.

6.2.2 Downstream OTR Images of Electron Beam

The next diagnostic we wish to analyze is the electron beam images taken on the downstream OTR screen. This screen is viewed by the "DTOTR2" camera, which looks at the back side of the OTR screen with a resolution of 30.3 μm/px and observes the transmitted light from the OTR emission process. Here, the imaging spectrometer images the horizontal electron beam size at an upstream object plane determined by the spectrometer quadrupole settings. [10] The vertical electron beam size is not imaged, and instead a dipole magnet upstream of the OTR screen disperses the electron beam vertically according to its energy. It is also important to keep in mind that electrons are only detectable if it produces enough OTR emission to be visible above the noise. If a slice of the beam is too disperse then it will not be visible.

propagates through more phase of the betatron oscillation process. This can result in an electron beam "focus" that is within the plasma.

[10] These settings assume a 10 GeV electron when calculating the object plane, so off-energy slices in the electron beam are slightly out-of-focus.

Table 6.4: Parameters of effective plasma lenses, assuming the Super-Gaussian fits of Fig. 6.3 correlate to the plasma density profile and the lens operates in the underdense regime. The 1 psi case assumes the same density profile shape as the 6 psi case.

Pressure (psi)	$\sim n_0$ ($\times 10^{16}$ cm^{-3})	L_{eff} (mm)	f (cm)	K (cm^{-2})	$\sqrt{K}L$
1	0.27	10.07	4.06	0.24	0.50
6	1.62	10.07	0.677	1.46	1.2
24	6.48	13.01	0.131	5.85	3.1
57.8	15.3	15.49	0.047	13.8	5.8
115.8	31.5	17.55	0.020	28.5	9.4

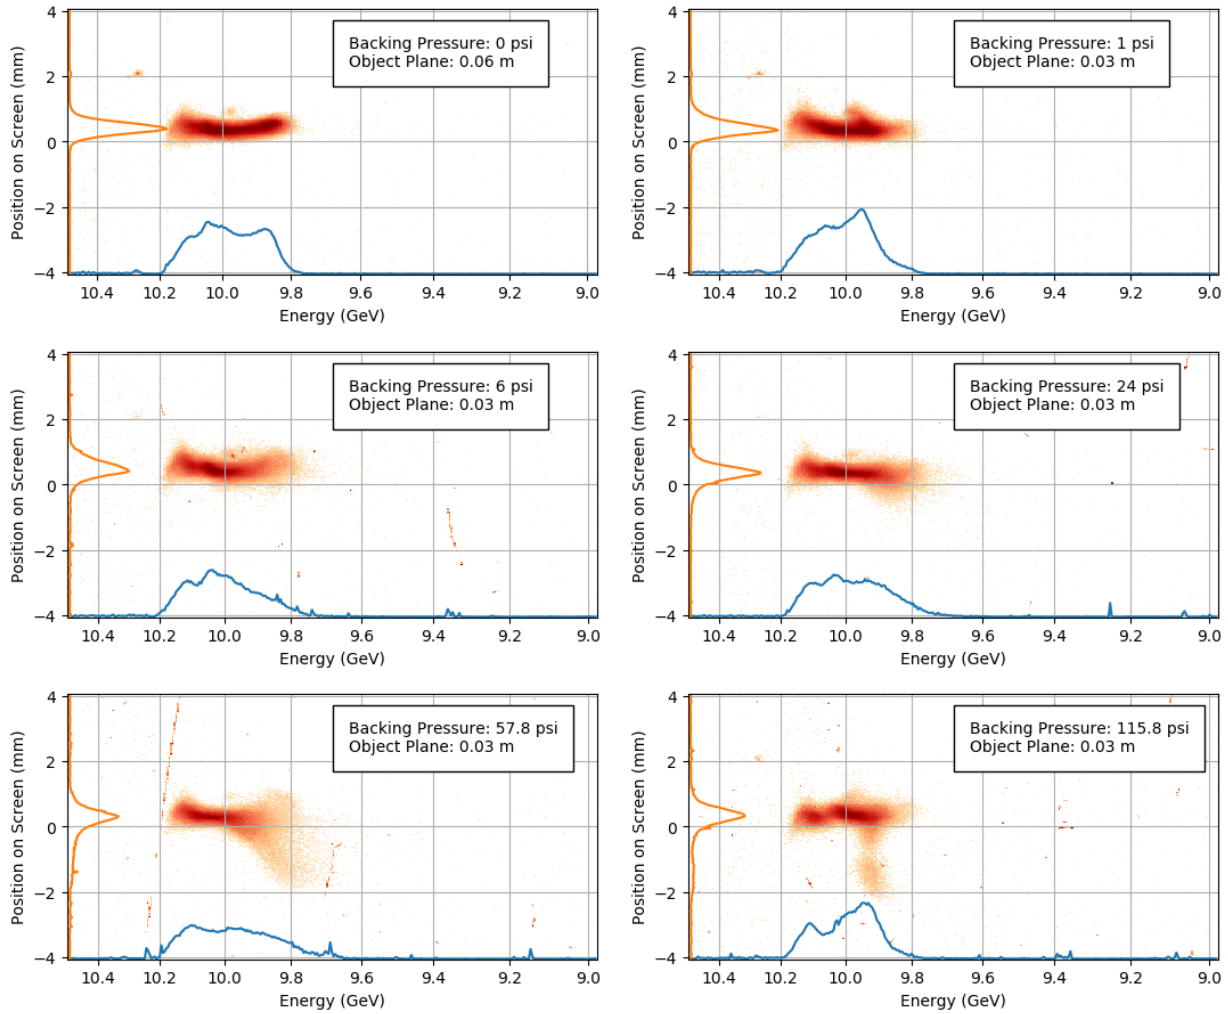

Figure 6.6: Sample images of the OTR light taken at the downstream DTOTR2 foil for various backing pressures. These images are taken with the object plane close to the position of the gas jet. The vertical axis on each subplot represents the imaged horizontal beam size at the object plane, and the horizontal axis are the electron energies after being dispersed by a dipole. The color scale is on a log scale, while the projections are on a linear scale. In orange on the left axis is a projection of beam size, and in blue on the bottom axis is a projection of energy distribution.

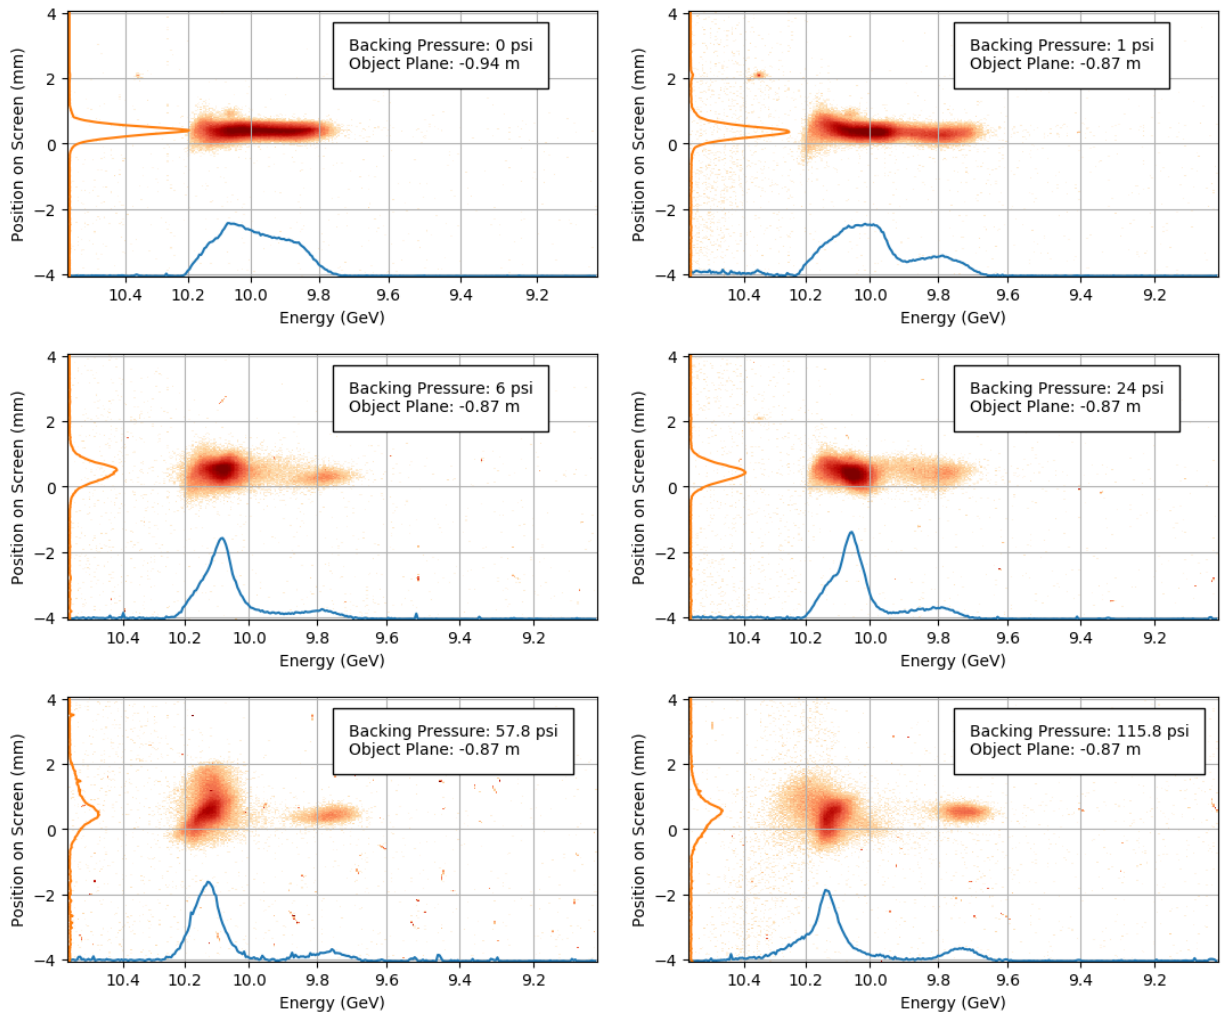

Figure 6.7: Sample images of the OTR light taken at the downstream DTOTR2 foil for various backing pressures. These images are taken with the object plane ~ 90 cm upstream of the gas jet. The vertical axis on each subplot represents the imaged horizontal beam size at the object plane, and the horizontal axis are the electron energies after being dispersed by a dipole. The color scale is on a log scale, while the projections are on a linear scale. In orange on the left axis is a projection of beam size, and in blue on the bottom axis is a projection of energy distribution.

Sample raw images of the light collected by the DTOTR2 camera are shown in Fig. 6.6 for an object plane close to the gas jet and in Fig. 6.7 for an object plane approximately 90 cm upstream of the gas jet. In these images, the light on the camera is plotted on a logarithmic scale while the projections onto either axis are plotted on a linear scale. The projection onto the vertical axis represents the beam's horizontal transverse size, and the projection onto the horizontal axis is a measure of the energy offset from the 10 GeV centroid.

Since the betafunction is on the order of ~ 1 m, we would expect the vacuum propagation to not vary too much on the order of ~ 1 m in drift space, which is what we see when comparing the the 0 psi dataset for these two positions. We start to see significant differences qualitatively between the two object plane locations as the backing pressure of the gas jet is increased. The beam's spectrum begins to show some distortions when the object plane is close to the gas jet. The closest we get to seeing a tight focus is the dispersed clouds of electrons between 9.9 and 9.8 GeV, which could show a tight focus quickly dispersing away from the focus. Further upstream we see evidence of what could be plasma wakefield acceleration. The single electron beam has a component that drives the wake and loses energy, and a component that is within the accelerating gradient and gains energy. However, it is interesting to note that this feature is only visible when the spectrometer is imaging an upstream object plane, as one would expect such a feature to appear in the energy spectra regardless of the object plane. One possible explanation could be within the exact imaging conditions for each energy slice of the beam, with some slices being out-of-focus and invisible relative to the noise.

To learn more, we need to examine these images more rigorously. First, we can examine the energy axis of the imaged beam using the horizontal projections. The previous plots of DTOTR2 use a calibrated energy scale to correctly plot the energy on the horizontal axis. This calibration is

$$E(x,z) = \frac{d_{nom} E_{bend}}{x + (d_{nom} - x_{beam}(z))}, \tag{6.1}$$

where $d_{nom} \approx 60$ mm is the nominal dispersion, $E_{bend} = 10$ GeV is the dipole setting, and $x_{beam}(z)$ is the position of the beam's energy centroid on the screen. This is found by plotting the energy

centroid in the 0 psi case (Fig. 6.8(a)). We make an observation here that the energy centroid on the DTOTR2 screen appears to shift when adjusting the imaging spectrometer settings, even for the case without a plasma. This is not a sign of the electron beam gaining or losing energy in the spectrometer, but it is a sign that the electron beam was initially not perfectly aligned going through the spectrometer. This will have further implications, but for now we fit the calibration $x_{beam}(z)$ as a linear function of the object plane z. In the calibrated comparison of Fig. 6.8(b) we see that the the presence of the plasma does impact the location of the energy centroid on DTOTR2. However, it is difficult to tell from this plot alone if this trend is due to the electron beam's energy or its vertical position and momentum (y,y') going into the spectrometer.

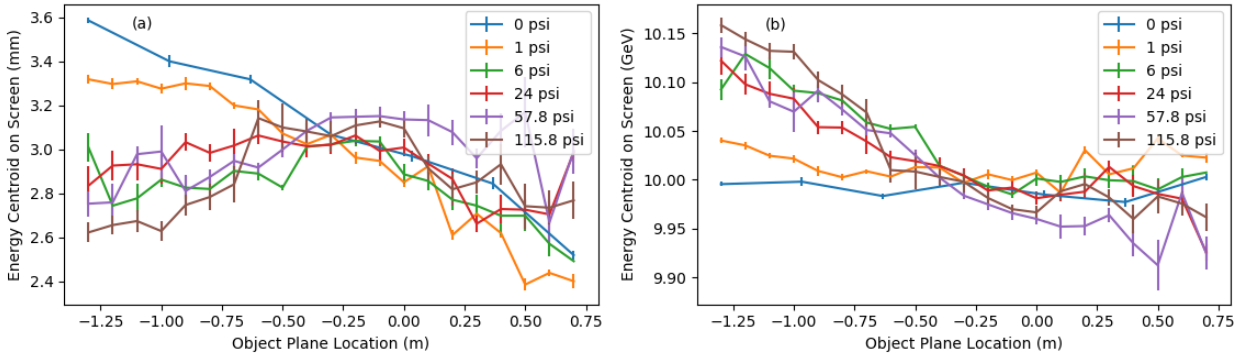

Figure 6.8: (a): Comparison of the position of the energy centroid across the scanned range of object plane positions for each backing pressure dataset. In particular, the energy centroid for the 0 psi case changes, although with no plasma in the experimental chamber the beam energy should still be 10 GeV. (b): From the positions with 0 psi where the beam centroid corresponds to 10 GeV, plotted here are the centroid's energy value across each dataset.

Next, we investigate the horizontal beam size by using projections of DTOTR2 onto the vertical axis. These projections of the beam size are fitted to a Gaussian, and we plot the evolution of this Gaussian spot size as the object plane is scanned in Fig. 6.9 for all datasets. This includes the magnification from the imaging spectrometer. At 0 psi the electron beam's measured size should agree with wirescanner measurements taken before and after the shift (Table 6.1), however the beam size measured using the wirescanner is $\sigma_{x,wire} = 38.6$ μm while the minimum beam size for the 0 psi object plane scan is $\sigma_{x,min} = 57.6$ μm. It is uncertain where the discrepancy arises from;

at this stage in commissioning either the wirescanners could produce an incorrect spot size, the spectrometer could not be imaging the desired object plane correctly,[11] or the DTOTR2 camera that views the OTR screen could be calibrated incorrectly. With this systematic uncertainty in horizontal beam size in mind, we continue with the goal of at least qualitatively measuring a change in the electron beam as the backing pressure is increased.

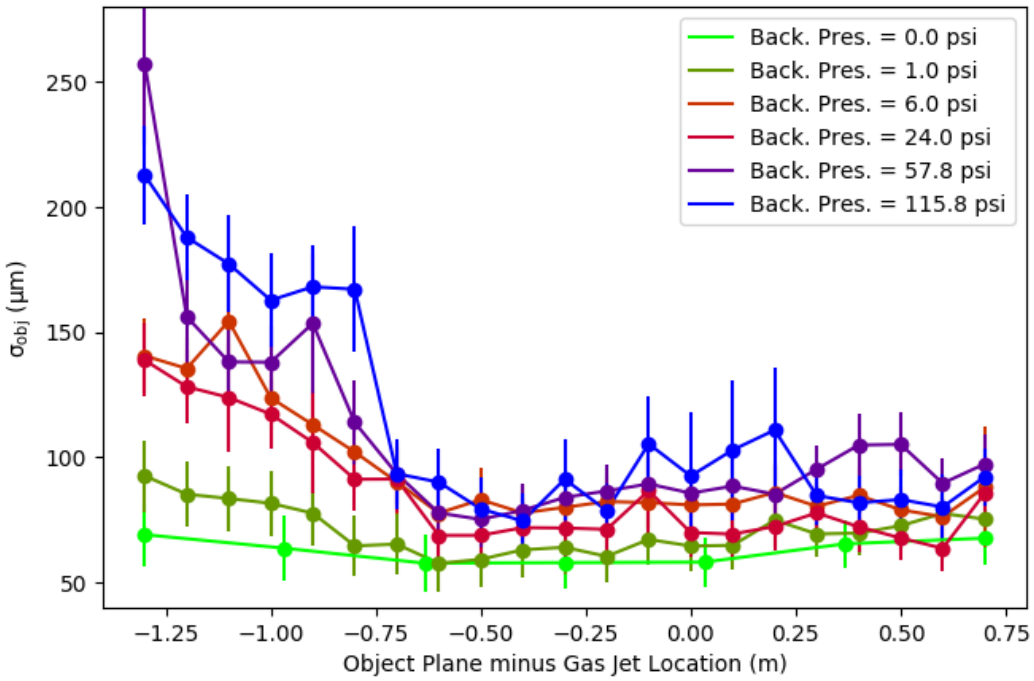

Figure 6.9: Projections of horizontal beam sizes on DTOTR2 plotted against the object plane location for increasing gas jet backing pressure. The 0 psi represents vacuum propagation with a $\sigma_{x,min} = 57.6$ μm focus. As the backing pressure is increased, the beam's divergence also increases as seen by > 70 cm upstream of the focus. Due to the single-bunch regime, it is difficult to resolve the focus itself because only a fraction of the beam undergoes focusing due to the plasma wake.

Further examining Fig. 6.9, as the backing pressure of the gas jet is increased the imaged spot size of the electron beam also generally increases. This is the opposite of what we would initially expect, but we also need to keep in mind that it is difficult for DTOTR2 to measure tightly focused beams. Rather, what we would see more clearly is that a beam with a smaller focus would diverge quicker when imaged away from the focus. We are also measuring the spot size by projecting across

[11] The beam not being aligned through the spectrometer doesn't help.

the full beam's spectrum, and in the single bunch regime we would expect only a fraction of the beam to actually be focused. A significant portion of the beam would be either unchanged or defocused from driving the wake interacting with the plasma electrons surrounding the wake. With this interpretation, the larger divergences for higher backing pressures in Fig. 6.9 when imaging far upstream of the focus seem reasonable.

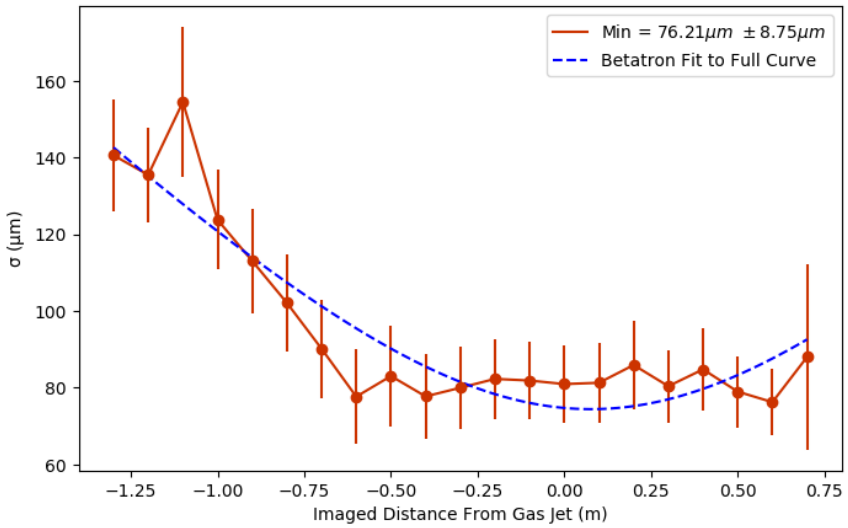

Figure 6.10: Sample of fitting a typical vacuum betatron evolution to the horizontal beam size measured on DTOTR2. Data in red is taken from the 6 psi dataset in Fig. 6.9, and the blue dashed curve is the fit to Eqn. 6.2.

To quantify this further, we can fit these spot size evolution plots to the standard equation for the beam's spot size evolution in vacuum drift space. Drift space propagation is easily modeled using CS parameters and transfer matrices,

$$\sigma(z) = \sqrt{\frac{\epsilon_N}{\gamma_L}\left(\beta^* + \frac{(z-z^*)^2}{\beta^*}\right)}, \tag{6.2}$$

and we use the normalized emittance ϵ_N, focused betafunction β^*, and focus location z^* as fit parameters. An example of such a fit is shown in Fig. 6.10 for the 6 psi dataset. We perform similar fits for all of the datasets and summarize the fit parameters in Table 6.5.

From this analysis of the full projection of the beam's horizontal size, we can note a few things. First, the 0 psi dataset does not match exactly with the values obtained before and after

Table 6.5: Betatron fit parameters using Eqn. 6.2 for the spot size evolutions using the full beam projections of Fig. 6.9. The spot size $\sigma_{min,measured}$ is the minimum horizontal size measured on DTOTR2 for each dataset, and is not a good measure of focusing strength in the single bunch regime.

Pressure (psi)	$\sim n_0$ ($\times 10^{16}\,\mathrm{cm}^{-3}$)	$\sigma_{min,measured}$ (μm)	β^* (cm)	z^* (cm)	ϵ_N ($\mu\mathrm{m}-\mathrm{rad}$)
0	0	57.6 ± 11.2	151.5	-0.30	42.8
1	0.27	57.5 ± 11.1	104.5	-0.20	71.0
6	1.62	76.2 ± 8.75	84.19	0.08	129
24	6.48	63.6 ± 8.95	80.12	0.09	107
57.8	15.3	75.1 ± 11.0	55.56	-0.03	191
115.8	31.5	74.5 ± 10.8	63.11	0.13	190

the shift using conventional diagnostics (Table 6.1). This can be due to a number of reasons as described earlier, and ultimately make it difficult to trust the exact numbers from this analysis of DTOTR2. However, looking at the general trends as the density increases, we do observe the focus β^* to decrease as the plasma density goes up. We also observe the emittance ϵ_N from the betatron fit increase with increasing density. This could be either due to chromatic phase spreading because the normalized plasma lengths $\sqrt{K}L$ are quite large (Table 6.4) and it could additionally be due to more collisions between the beam and plasma particles as the density increases. The electron beam's transverse size is also quite large so it is possible for the wake to be smaller than the beam, which would also lead to aberrations in the focusing force and an increase in emittance. We will return to these fit parameters in Sec. 6.3 when we compare with theoretical predictions.

Lastly, we can attempt to dig further into this DTOTR2 data by using horizontal projections for different energy slices in the beam. While we do not know what the original longitudinal phase space of the beam was during the commissioning shift, it is possible for an electron beam to possess a "chirp." This would mean that the electron has a correlation between its longitudinal position ξ and its energy δ. If this were to be the case, then for a single-bunch plasma lens experiment we would see different effects for high and low energy slices. It would then be possible to separate the beam sizes for slices that are focused by the plasma wake from slices that are only responsible for driving the wake. Again, there is no indication that this was the case for the electron beam during the shift, but regardless we can perform this analysis and see if there are signs of a chirped electron

beam.

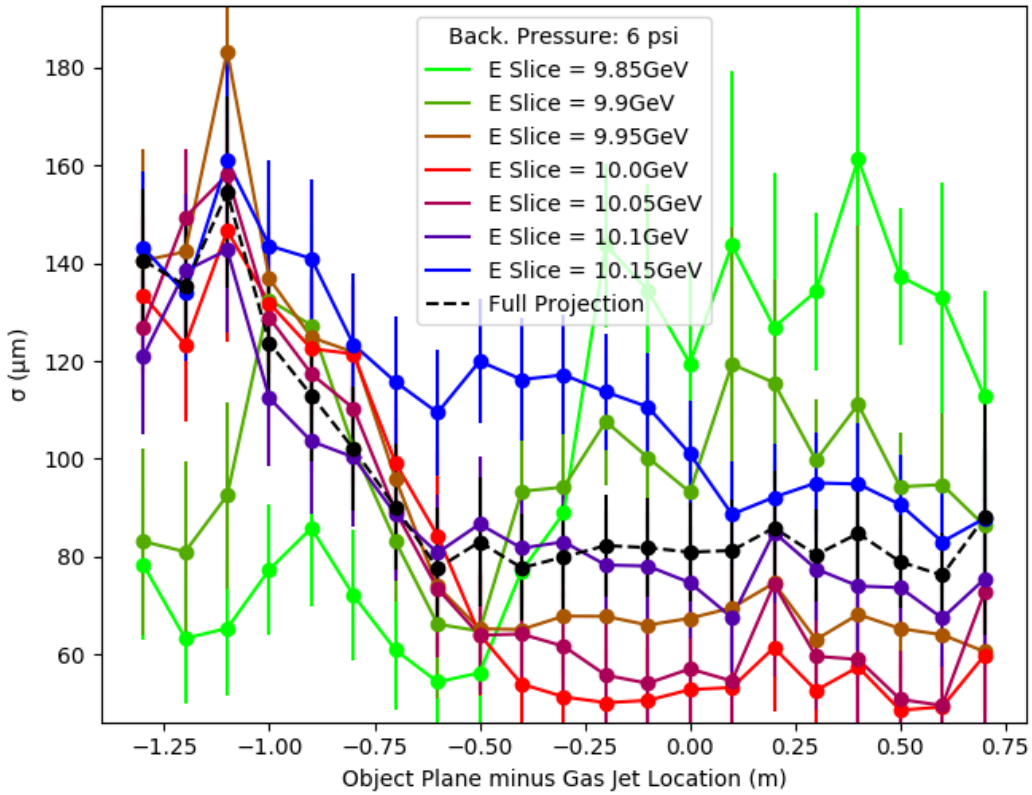

Figure 6.11: Evolution of energy projection slices' horizontal beam size as imaged on DTOTR2. Horizontal axis is the object plane from the spectrometer magnet settings, and the vertical axis is the beam size including the 30.3 μm/pix resolution and magnification from the spectrometer settings. The different colors represent the various energy slices the sigmas are evaluated at, with the errorbars coming from statistical errors of 10 shots at each step and the systematic error of the DTOTR2 camera resolution. Dashed black is the spot size evolution for the beam's full projection. Each energy slice represents a bin of electron energies within $\pm$0.025 GeV, and the slice position on DTOTR2 is determined by the 10 GeV centroid position with 0 psi from Fig. 6.8.

The procedure is as follows: First we take the DTOTR2 images and re-plot the axes with the energy scale. Then, we divide up the beam spectrum into bins of energy slices in 0.05 GeV intervals. In each bin, we project the camera signal onto the horizontal beam size axis and fit the OTR light to a Gaussian spot size. In Fig. 6.11 is an example of the spot size evolution for the various energy slices with a backing pressure of 6 psi, compared with the spot size evolution for the beam's full projection. We do see that there are beam slices that have a smaller spot size than the full projection, such as the 10 GeV slice. However, before assuming that this is evidence of a

chirped beam, we also note this could be either from the 10 GeV energy slice being imaged more correctly or from imaging aberrations in the misaligned spectrometer.

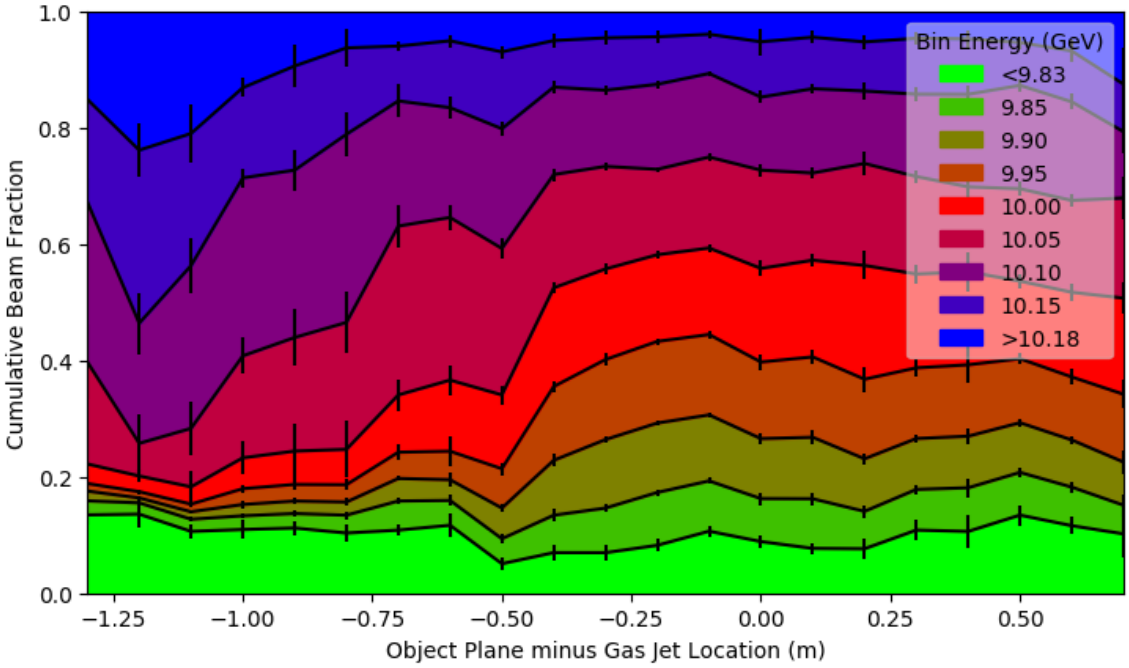

Figure 6.12: Cumulative fractions of the electron beam in different bins of energy for the 6 psi dataset. Horizontal axis is the object plane location, and vertical axis plots the relative percentage of the beam observed in a region of the energy projection. When imaged far upstream of the gas jet location, more signal is visible from where higher energy electrons would be on the DTOTR2 screen. Each central bin is a region of ± 0.025 GeV around the stated energy in the legend.

The distribution of beam electrons with energy also changes with the object plane, as plotted in Fig. 6.12 for the 6 psi dataset. A similar plot for the 0 psi dataset shows a constant representation across the object plane scan, indicating that this is an effect of the plasma interaction. In Sec. 6.2.3 we will see the total charge on DTOTR2 only has minor variation across the object plane scan, so this change in relative distribution is not due to charge loss. An explanation for the relative distribution change is certain energy slices being out of focus, but then one would expect the total signal to significantly decrease. Once again, this is most likely a combination of a misaligned spectrometer with energy gain and/or loss in the plasma interaction.

Next, we take the spot size evolutions for all of the energy slices across all datasets of rising

backing pressure and perform betatron fits using Eqn. 6.2. In Sec. 6.3 these fit parameters are plotted in Figs. 6.15 and 6.17, and they are compared with the fit parameters of the full projection and with theoretical predictions. Most of the curves produce a decent fit, although the high energy slice at 0 psi is too weak and cut from the figures. Overall trends of decreasing beta and increasing emittance agree with what we found for the betatron fits of the full projection, but there is little evidence that specific energy slices behaved significantly different from the full beam. Therefore, we do not think that the electron beam from this commissioning shift had a clear energy chirp.

Another aspect that could be at play is energy change in the beam. For an ideal plasma lens with $\sqrt{K}L \ll 1$, the energy gain is negligible and this is not an issue. The plasma from this commissioning shift, however, are much larger (Table 6.4). It is possible that different longitudinal slices of the beam undergo different accelerating and decelerating fields depending on the length of the plasma wake. This further muddies the water when attempting to separate out constant longitudinal slices of the beam across datasets with different plasma densities.

6.2.3 Charge Measurements

We can make some observations on the electron beam charge using both toroid measurements of beam charge[12] and the amount of light collected on DTOTR2. We first plot the toroid readings along the linac during each dataset in Fig. 6.13. The toroids that were far upstream of the experimental area (red, green, orange) all agree that the incoming electron beam's charge was constant with little variation between shots. The toroid represented by the brown curve is 1.35 m upstream of the gas jet and only shows a varying signal at the upper limit of the scanned gas jet backing pressures. In grey, the toroid 2.72 m downstream of the gas jet shows large measurement fluctuations and a much higher charge reading. This increase in measured charge is more likely due to the presence of scattered plasma electrons interfering with the diagnostic rather than the electron beam gaining a significant amount of charge.

For a clearer measurement of the amount of charge in the electron beam after the plasma

[12] Toroids generate a voltage proportional to the electron beam current that passes through it.

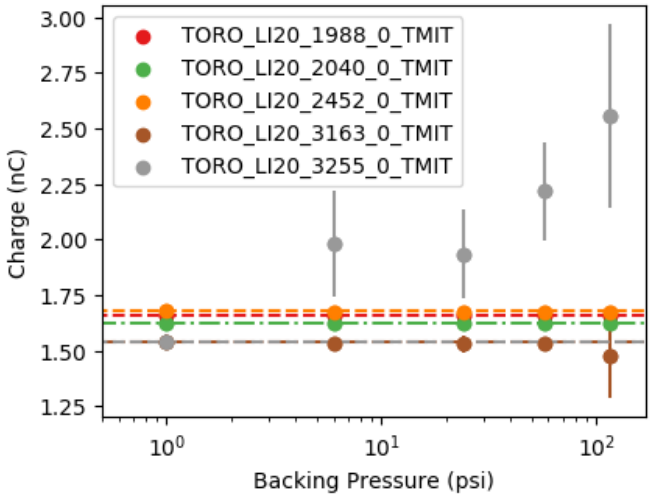

Figure 6.13: Toroid signals along the linac which measure the number of electrons passing through every shot. The horizontal axis is the gas jet backing pressure for the given dataset. The red, green, and orange curves represent toroids located far upstream of the experimental chamber in the linac, and are unaffected by the operation of the ionized gas jet. The brown curve is a toroid located 1.35 m upstream of the gas jet, and the grey curve is a toroid located 2.72 m downstream of the gas jet. The dashed curves represent that toroid's average signal at 0 psi backing pressure. Differences between toroid signals are due to the devices being uncalibrated at the time.

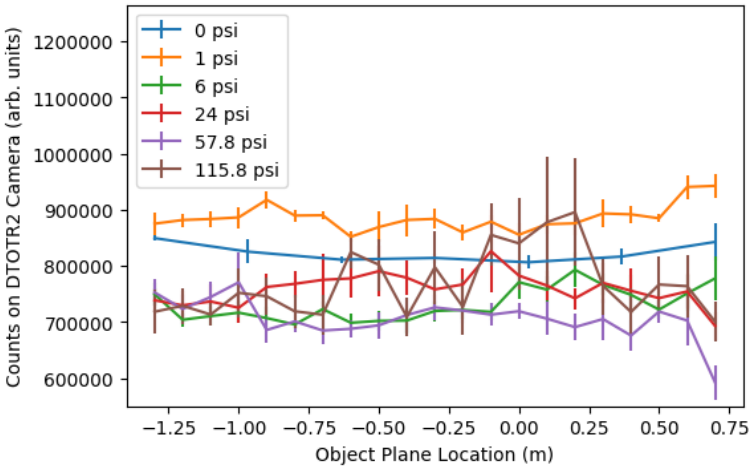

Figure 6.14: Comparison of the total amount of visible OTR light on DTOTR2 for each backing pressure dataset across the full range of scanned object planes. Slight variation of visible light at the 0 psi dataset suggests imaging conditions affect intensity visible on camera. Images with a backing pressure above 6 psi have overall lower intensity on the camera, which implies the electron beam experienced some charge loss.

interaction, we return to the DTOTR2 images. The visible OTR light above the noise on the screen

can be correlated to the charge present in the electron beam as it interacts with the OTR foil. This is plotted in Fig. 6.14 for all of the datasets across the object plane positions. Looking at the 0 psi dataset, the is only a slight variation of DTOTR2 camera intensity as the object plane is scanned which may indicate a small amount of charge loss in the spectrometer. However, it is more likely just a result of the imaging conditions since it would be much more drastic if the electron beam was hitting something in the spectrometer.

Comparing the other datasets of Fig. 6.14, the 1 psi dataset has slightly more visible light off of DTOTR2 while the datasets a higher backing pressures have slightly less visible light. At higher backing pressures, the electron beam could lose some charge from either plasma collisions or simply being strongly deflected by transverse fields in the plasma. If electrons have too large of a divergence, then they will either be too weak at DTOTR2 to produce visible radiation or will be lost to the beam-pipe walls all together. The increase in signal at 1 psi is less understood. Electrons from the plasma would not be able to make it to DTOTR2 because their longitudinal energies are much less than 10 GeV and the spectrometer magnets would completely deflect them away. The signal increase, in one way or another, is due to more of the beam electrons being visible above the noise on DTOTR2.

6.3 Comparison to Theory

It is not straightforward to compare the results of this commissioning shift to the theory of ideal plasma lens operation presented in Chapter 3. We are in the single-bunch regime, where only a fraction of the beam would witness the strong focusing force to begin with. The plasma itself is quite long, and from Table 6.4 we see that only at 1 psi are we in a regime where the plasma can be considered a "lens" with $\sqrt{K}L < 1$. A longer plasma also means that we will have to contend with energy change in the beam. And, to top it all off, the estimated electron beam density from Table 6.1 suggests we are only in the nonlinear blowout regime with $n_{beam} > n_{plasma}$ up until 6 psi backing pressure. While most of the datasets exist in invalid regimes, the 1 psi dataset has the best possible chance of being in the underdense plasma lens regime.

Despite this, we can still make an attempt to compare these commissioning results with theory as an exercise to see (1) if we do understand the regime boundaries correctly and (2) if we do observe underdense plasma lens focusing for the 1 psi case. First, we plot the waist betafunctions with the increasing estimated plasma density in Fig. 6.15. At $n_p \approx n_{beam} \approx 1.7 \times 10^{16}$ cm^{-3}, the plasma-beam dynamics are expected to shift from the nonlinear blowout regime to the linear regime, and this boundary is marked in dotted black. The fit parameters for the energy slices (yellow to red) are not too different from the fit parameters for the full projection (black). Some slices appear more focused than others, but this can be an effect of the off-energy slices not being imaged as well.

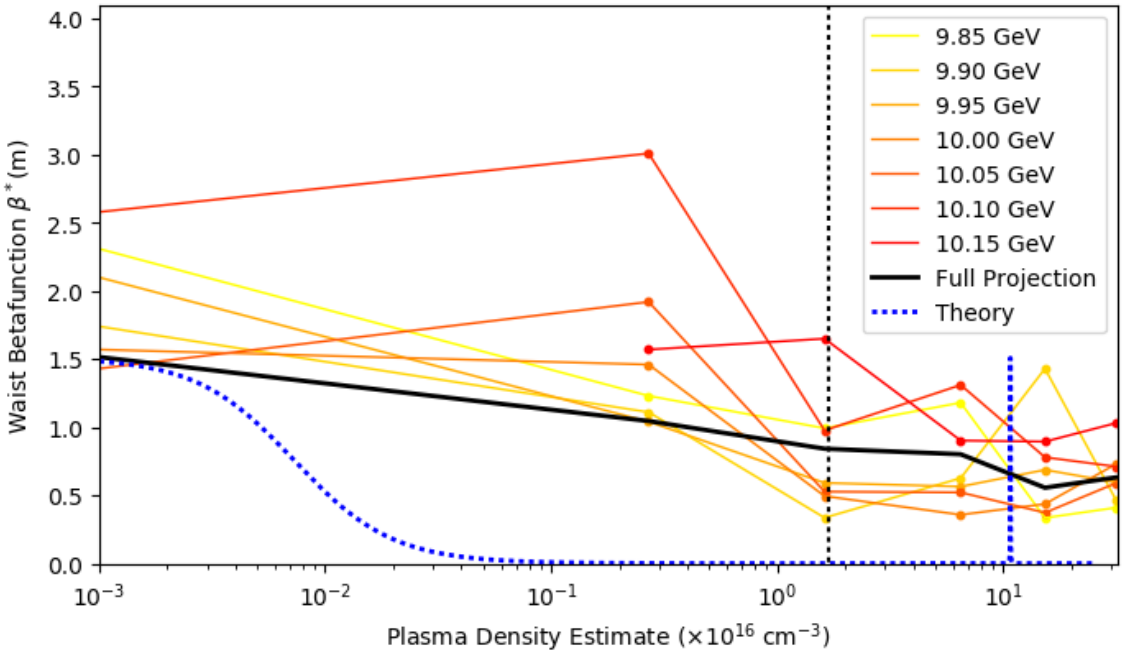

Figure 6.15: Comparison of betafunction fit parameters for the full projection (black), energy slice projections (yellow to red), and the theoretical predicted betafunction focus from ideal underdense focusing (dotted blue). The vertical dotted black line is the approximate transition from the nonlinear regime (left) to the linear regime (right). While the energy slices roughly track the betafunction of the full projection, both are significantly above the theoretical curve. The location where the theory curve is vertical represents where the normalized thickness $\sqrt{K}L = 2\pi$ is such that the beam undergoes a full betatron phase-space rotation in the plasma.

More significantly, the theoretical prediction for the waist betafunction is orders of magnitude lower than what we observe. For the 1 psi case, if we assume ideal nonlinear focusing for an initial

betafunction of $\beta_i = 151$ cm, a plasma length of 10.1 mm, and a density of $n_p = 2.7 \times 10^{15}$ cm^{-3}, we would expect to see an impressive betafunction focus of $\beta^* = 1.2$ mm (Fig 6.16). Since we are in the single-bunch regime only a fraction of the beam would be focused to this degree, and since there is little evidence for a chirped beam we cannot disentangle this beam's component with only the spectrum on DTOTR2. It is also unlikely we would resolve such a tight focus with a betafunction of ~ 1 mm, as the step sizes used in the data collection are a few cm's long. Furthermore, if there were to be a fraction of the beam focused this strongly, it may have some difficulty in clearing some of the tighter apertures in the spectrometer. With a 1.2 mm betafunction focus 4 cm downstream of the gas jet, this such a beam would have a transverse spot size of 4.1 mm at the location of the downstream Beryllium window. This window has a radius of 5 mm, and so it is feasible that highly-divergent electrons are being lost at this aperture.

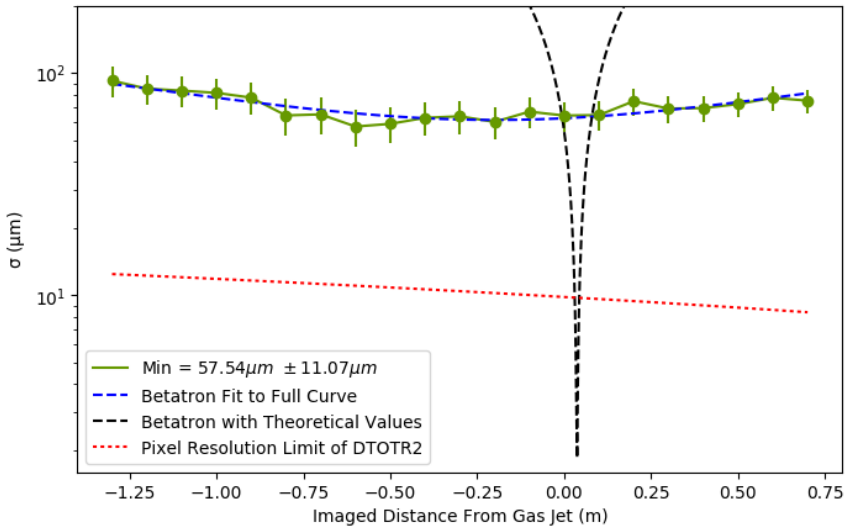

Figure 6.16: At the 1 psi backing pressure dataset, plots of the horizontal beam size as measured by DTOTR2 (green) along with the respective betatron fit from Table 6.5 (blue-dashed). The theoretical prediction for ideal nonlinear focusing is plotted in dashed black. The ultimate lower limit given by DTOTR2's pixel resolution is plotted in dotted red.

Next we plot the emittance fit parameters in Fig. 6.17 with a similar plotting style marking the linear-nonlinear boundary and comparing the energy slice projections with the full projection. When comparing to theoretical predictions of the emittance growth assuming ideal nonlinear regime

operation, here the emittance values agree fairly well to theory up until the linear regime. In fact, if we only assume chromatic phase spreading one would expect the beam's full projected emittance to grow while the emittances of the separate energy slices remain approximately constant. This is seen to some degree in Fig. 6.17 when looking at the individual energy slice emittances. Once at the linear regime, the transverse wakefields are nonlinear functions in radius r and depend on the drive beam density. We would expect the emittance growth be significant and to stop scaling as strongly with plasma density, which Fig. 6.17 also suggests. However, increasing plasma density can lead to emittance growth through Coulomb scatter off of the plasma electrons and ions, and this effect roughly scales linearly with the plasma density [107].

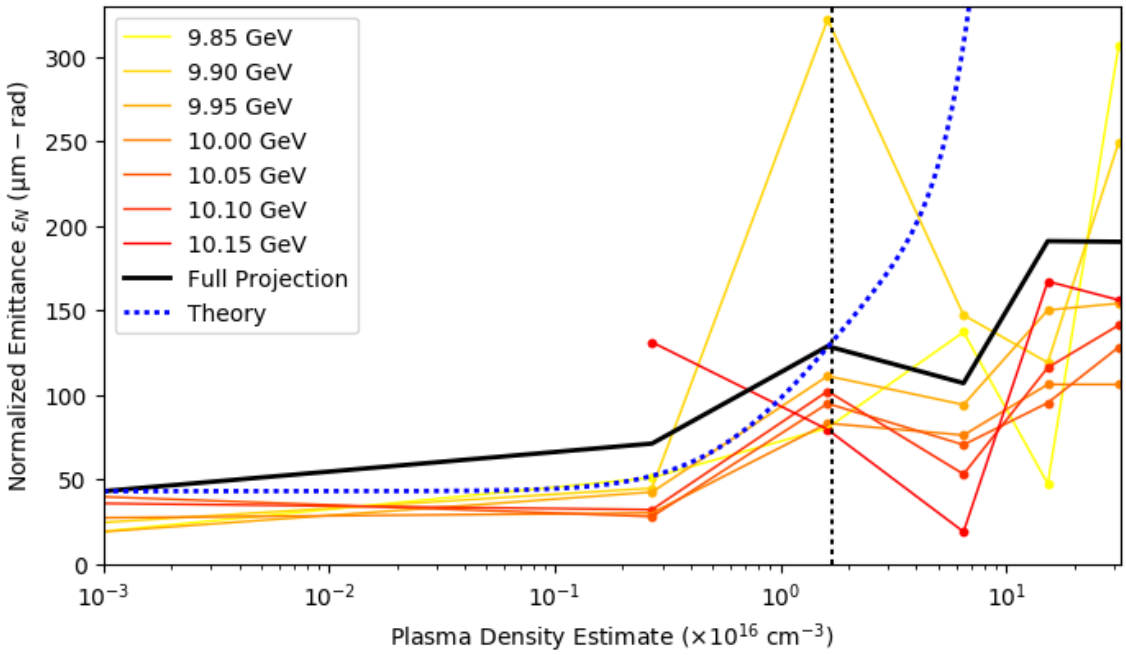

Figure 6.17: Comparison of the normalized emittance fit parameters for the full projection (black), energy slice projections (yellow to red), and the theoretical emittance growth due to chromatic phase spreading in the underdense regime (dotted blue). The vertical dotted black line is the approximate transition from the nonlinear regime (left) to the linear regime (right). In the nonlinear regime, the full projection's emittance increases more than the emittance of individual energy slices, which is expected. Then, in the linear regime, the emittance values grow considerably more due to the nonlinear focusing force and possible Coulomb scattering contributions.

Lastly, we wish to give a brief note on the effects energy change can have on this data. Much

of the analysis assumes that the energy is constant before and after the plasma interaction. This is true for a plasma lens, but here the plasma lengths are significant (Table 6.4) and the longitudinal phase space of the beam will evolve. It is difficult to predict the full longitudinal wakefields across the length of the beam for the single-bunch regime, especially as we vary the density from the nonlinear to linear regimes, but we can look at rough scaling equations for the peak longitudinal wakefields.

In the nonlinear regime, the peak accelerating gradients can be estimated with the following equation from Ref. [39]:

$$eE_{max} \approx 1\frac{\text{eV}}{\text{cm}} \times \sqrt{n_p[\text{cm}^{-3}]}. \tag{6.3}$$

For the 1 psi case with $n_p \approx 2.7\times10^{15}$ cm^{-3}, Eqn. 6.3 gives a possible energy increase of 0.05 GeV in 1 cm of plasma if there are beam electrons in this optimal phase of the wake. A similar decelerating field can be experienced in the front half of the wake where the drive beam is located. This amount of energy gain and/or loss can be enough to move electrons between the different energy slice bins used in the analysis of Sec. 6.2.2. As we move into the linear regime, the wakefields start to scale with the beam density rather than the plasma density, and so we can expect peak longitudinal wakefields on the order of $\pm$0.13 GeV/cm. This can be enough to significantly alter the longitudinal phase space of the electron beam for the higher backing pressure datasets. Depending on the exact longitudinal coordinate of each electron and the dimensions of the plasma wake, electrons can lose or gain up to this significant amount of energy. As a result, it is difficult to compare energy slices of the beam across the various datasets for this commissioning shift. On the other hand, these rough calculations do support the evidence of small-scale plasma wakefield acceleration seen in the raw DTOTR2 images of Fig. 6.7.

Chapter 7

Discussions

In this dissertation we have introduced the underlying physics behind passive underdense plasma lenses, investigated the transverse focusing quality from a theoretical perspective, and reviewed design considerations in pursuit of the goal of generating a laser-ionized passive plasma lens. We have made experimental progress towards demonstrating underdense plasma lensing at the FACET-II facility, and the experimental plasma lens program there is poised for attaining results in the near future. In Sec. 7.1 we review the major conclusions from this work, and in Sec. 7.2 we discuss future experimental goals of the plasma lens program at FACET-II and the prospects of underdense plasma lenses in the greater accelerator community.

7.1 Conclusions

Using the analytic framework of electron beams in the nonlinear blowout regime of PWFA, we are able to model passive, underdense plasma lenses with transfer matrices. This not only allows for the rms spot size focus to be easily calculated for both thin and thick regime plasma lenses, but we can also calculate the chromatic emittance growth in these plasma lenses. We found analytic expressions for the emittance growth in a plasma lens, and gained more confidence that in the thin regime these lenses do not contribute too much to emittance growth.

We introduced various hypothetical applications for a plasma lens using this analytic description of the focusing dynamics. One such application is that of matching a witness bunch into a plasma wakefield accelerator for emittance preservation. In order for a beam to avoid chromatic

phase spreading in a long PWFA stage, that beam must have a very small and specific beam size at the start of the plasma stage. This leads to a matching condition, and this condition must be satisfied using upstream focusing tools. The matched spot size is generally too small to be within the reach of conventional quadrupole electromagnets, and plasma density ramps that might provide sufficient focusing are challenging to realize in practice. On the other hand, we have shown that plasma lenses can easily achieve these matching conditions, and we have also derived expressions to calculate the required plasma lens location and thickness if the incoming beam parameters and matching condition is known.

Additionally, we performed some theoretical studies on applying underdense plasma lenses as a strong final focus and the resulting focusing aberrations due to synchrotron radiation. As part of this we found a hypothetical parameter space where the Oide limit could be reached for the first time experimentally using an underdense plasma lens at FACET-II. We also used numerical PIC simulations to study underdense plasma lenses in the ideal two-bunch configuration, as well as more exotic configurations more difficult to analytically model. These include the use of the second blowout wake period for focuisng the trailing witness bunch, and single bunch focusing, where a single electron beam drives a wake with the front while the bulk of the beam experiences the strong focusing. These studies show that the underdense plasma lens formalism can be partially extended to these extreme case studies, and give some motivation to future experimental efforts.

With the performance of the passive underdense plasma lens understood, we moved on towards designing an experimental setup to generate these plasma lenses with the required dimensions. Laser ionization with a transversely-propagating laser of moderate intensity was chosen as a suitable process to ionize a plasma volume with longitudinal thicknesses on the order of $10's$-$100's$ μm. To that end, two optical setups are presented that give fitting plasma density profiles: one with crossed cylindrical lenses and one with a single spherical lens. The effects of laser refraction due to the ionization process is simulated in both a static, uniform volume of gas and in the outflow of a gas jet. It was found in that the laser can still produce the desired plasma volume despite these effects.

Furthermore, on the topic of gas jets we presented a detailed analysis of the effect a transverse plasma density gradient has on the wakefields of the nonlinear blowout regime. Such density gradients can be found in the outflow of gas jets, and these gas jets are used in a variety of plasma-based acceleration and focusing devices beyond even just the underdense plasma lens. The transverse wakefields due to the nonuniform ion column is found analytically, and we are able to use an empirical model to estimate the contribution from the non-axisymmetric electron sheath. We used this model to predict the net deflection on a witness bunch and the agreement was very strong. In the future, this model can be applied to estimate the permissible tolerances in the design of a plasma lens for applications with particularly strict requirements on transverse deflection, as well as prescribe target parameters to minimize harmful aspects of these wakefields. Alternatively, one could also try to purposefully design a plasma lens with a large density gradient to make a novel plasma-based beam kicker.

Using our understanding of the laser ionization of plasma lenses, we proposed an experiment at FACET-II to study these passive plasma lenses with 10 GeV electron beams. Our setup involves a laser propagating transverse to the electron beam axis that ionizes a gas jet outflow using a low energy probe arm. While we have made progress towards generating such a plasma lens in the accelerator tunnel, work still needs to be done to properly ionizing a plasma using this setup. In particular, the laser profile needs to be optimized and there needs to be more alignment controls installed to combat the complexity of the FACET-II probe setup, and the B-integral for the ionizing laser needs to be addressed further through new solutions at reducing the amount of non-vacuum material the laser propagates through.

Despite the side-ionization laser not being fully ready, we were able to commission some aspects of the plasma lens experiment using the longitudinally-propagating, high-power main laser line. In a commissioning shift, we ionized an elongated gas jet using the focus of an axilens optic, which resulted in a fairly long plasma of 1-2 cm. The current FACET-II beam parameters at the time were a single electron bunch with a fairly large transverse spot size and low density. While this shift did not have our intended, ideal parameters of a two-bunch electron beam with a small initial

spot size propagating through a thin plasma, we were able to test the experimental procedures and data acquisition system, as well as develop analysis tools and techniques for studying the plasma lens. There was one dataset that had the potential to be in the passive underdense plasma regime, but there was unfortunately not enough evidence to support that underdense plasma lensing had occurred. If the longitudinally-ionized plasma and single electron bunch were to be used again for a plasma lens experimental shift, then we now know the approximate signals we are expecting and can fine-tune the experiment to have a smaller longitudinal step size, a more narrow gas jet instead of the elongated nozzle, and an electron beam with a smaller initial spot size. We are confident that in the future we will successfully observe underdense plasma lensing once these small improvements are made to the experiment.

7.2 Future Work

The experimental efforts for the underdense plasma lens, and even those related to our design for the FACET-II plasma lens, are far from finished. The commissioning shift was a good first step in gaining the necessary experience required for future shifts, but the parameter space and setup was far from ideal. On the electron beam front, as FACET-II final installation processes reach their conclusions the parameter space available to the electron beam will begin to resemble the designed values more and more. This not only means access to smaller beam sizes that would behave more subtly to the transverse focusing forces of a blowout plasma wake, but also access to a second electron bunch to do proper passive plasma lensing with a two-bunch configuration.

On the laser front, building up the FACET-II probeline suite further will eventually allow for the intended operation of the transversely propagating ionization probe laser. Over the past few months already has enormous progress been made in increasing the quality of FACET-II's laser, and so we are confident that the system will be more stable and understood as operation and alignment procedures become mature. Improvements will be made towards maximizing the laser quality in the various probe arms, and this includes procedures and devices to help align the out-of-vacuum optics with higher precision than what is currently available. Once the ionizer probe

is up and running, we will be able to run plasma lens experiments using a very thin plasma with the additional capacity of being able to scan the longitudinal thickness by shifting the laser focus. This will be of enormous value in future plasma lens experiments to demonstrate that the focal length can be adjusted to high precision.

Once these initial plasma lens experiments using the intended setup are accomplished, then the plasma lens at FACET-II will be in a strategic position to assist with the emittance preservation studies for PWFA. This research area of PWFA is one of the foremost goals of the FACET-II facility, with the nominal location of these PWFA plasma sources located just downstream of where the plasma lens is currently installed. As we've demonstrated in numerical particle tracking simulations, depending on the initial parameters and the matching conditions a plasma lens of very moderate thickness and density can be enough to achieve matching. One can easily imagine an experiment where a PWFA acceleration stage is utilized with and without a correctly-designed upstream plasma lens to compare the differing degrees of emittance preservation.

Aside from PWFA Matching, the parameter regime of FACET-II can allow for further scientific exploration using underdense plasma lenses. A setup which uses a very dense plasma lens with a very wide witness beam is able to generate enough synchrotron radiation that the beam's focus becomes significantly distorted. This Oide limit has not yet been reached experimentally, and we have shown that the plasma lens has a parameter space that can allow for this limit to be reached. In fact, this parameter space is not far off of what FACET-II is currently aiming towards. To help reach this parameter space, one could also explore the regime of second wake period focusing, which we have performed numerical studies on.

Beyond the plasma lens experiment of FACET-II, there is a need for strong axisymmetric focusing in a variety of accelerator applications that underdense plasma lenses are candidates in fulfilling. In both PWFA and LWFA experiments, at the exit of the plasma stage electron beams often have extremely high divergences and it is difficult to couple the beams into other accelerator components downstream. The high focusing strengths and small longitudinal footprint of the underdense plasma lens can be applied towards controlling the divergence of these beams in a short

distance. This can be extremely beneficial in the case of LWFA experiments that generate the accelerated beam in the plasma and wish to use the beam for experiments without needing a long accelerator hall, such as in an x-ray free electron laser (FEL).

This general application of using plasma lenses for staging can be applied towards large accelerator designs as well. While most of the current, ongoing experimental efforts around PWFA and LWFA involve optimizing a single acceleration stage, the challenge now looms in the community to demonstrate that a witness beam can be coupled out of one accelerator stage and into another without degrading the beam quality. Doing this in a small longitudinal footprint is also ideal, as it will allow for next-generation accelerators to optimize their cost-to-performance ratio. In such an environment, one can imagine a standardized underdense plasma lens to be an important component in providing necessary strong focusing fields.

Lastly, it would be improper to end this section without discussing the prospects of using underdense plasma lenses as a final focusing optic for high energy colliders. Future work must certainly be undertaken to investigate the suitability of underdense plasma lenses for colliders, but there are a few pros and cons we can already point to. The main proponent of underdense plasma lenses for this application are the already discussed focusing fields themselves, which are several orders of magnitude stronger than other options. With the goal of increasing luminosity through smaller spot sizes, this is already a great start to the underdense plasma lens resume. However, there are three big concerns that need to be addressed. First is the aforementioned Oide limit, which may pose a restriction on just how far the underdense plasma lens can go. Second, the underdense plasma lens is only capable of focusing the negatively charged electrons in the blowout regime. If one wants to focus a positron beam then the plasma and blowout wake need to be very carefully designed, and even so there is no guarantee on the quality of these focused positron beams. Lastly, the impact of having a plasma lens close to high-precision detector equipment will need to be evaluated. If a simple gas jet plasma lens is used in this environment, it is easy to imagine the neutral gas or scattered plasma particles interfering with collision diagnostics.

Bibliography

[1] J W Wang and G A Loew. Field Emission and RF Breakdown in High-Gradient Room-Temperature Linac Structures. Technical Report SLAC-PUB-7684, Stanford Linear Accelerator Center, Stanford Univ., CA, 1997.

[2] Nan Phinney, Nobukasu Toge, and Nicholas Walker. ILC Reference Design Report Volume 3 - Accelerator, December 2007. arXiv:0712.2361 [physics].

[3] T. Tajima and J. M. Dawson. Laser Electron Accelerator. Phys. Rev. Lett., 43(4):267–270, July 1979. Publisher: American Physical Society.

[4] R. D. Ruth, A. W. Chao, P. L. Morton, and P. B. Wilson. A Plasma Wake Field Accelerator. Particle Accelerators, 17(6):171, 1985.

[5] Pisin Chen, J. M. Dawson, Robert W. Huff, and T. Katsouleas. Acceleration of Electrons by the Interaction of a Bunched Electron Beam with a Plasma. Phys. Rev. Lett., 54(7):693–696, February 1985. Publisher: American Physical Society.

[6] J. B. Rosenzweig, B. Breizman, T. Katsouleas, and J. J. Su. Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields. Phys. Rev. A, 44(10):R6189–R6192, November 1991. Publisher: American Physical Society.

[7] Ian Blumenfeld, Christopher E. Clayton, Franz-Josef Decker, Mark J. Hogan, Chengkun Huang, Rasmus Ischebeck, Richard Iverson, Chandrashekhar Joshi, Thomas Katsouleas, Neil Kirby, Wei Lu, Kenneth A. Marsh, Warren B. Mori, Patric Muggli, Erdem Oz, Robert H. Siemann, Dieter Walz, and Miaomiao Zhou. Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator. Nature, 445(7129):741–744, February 2007. Number: 7129 Publisher: Nature Publishing Group.

[8] M. J. Hogan, T. O. Raubenheimer, A. Seryi, P. Muggli, T. Katsouleas, C. Huang, W. Lu, W. An, K. A. Marsh, W. B. Mori, C. E. Clayton, and C. Joshi. Plasma wakefield acceleration experiments at FACET. New J. Phys., 12(5):055030, May 2010.

[9] M. Litos, E. Adli, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. P. Delahaye, R. J. England, A. S. Fisher, J. Frederico, S. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, N. Vafaei-Najafabadi, D. Walz, G. White, Z. Wu, V. Yakimenko, and G. Yocky. High-efficiency acceleration of an electron beam in a plasma wakefield accelerator. Nature, 515(7525):92–95, November 2014. Number: 7525 Publisher: Nature Publishing Group.

[10] M. Litos, E. Adli, J. M. Allen, W. An, C. I. Clarke, S. Corde, C. E. Clayton, J. Frederico, S. J. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, M. Schmeltz, N. Vafaei-Najafabadi, and V. Yakimenko. 9 GeV energy gain in a beam-driven plasma wakefield accelerator. Plasma Phys. Control. Fusion, 58(3):034017, February 2016. Publisher: IOP Publishing.

[11] Technical Design Report for the FACET-II Project at SLAC National Accelerator Laboratory. Technical Report SLAC-R-1072, SLAC National Accelerator Laboratory, CA, 2016.

[12] W. P. Leemans, B. Nagler, A. J. Gonsalves, Cs Tóth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker. GeV electron beams from a centimetre-scale accelerator. Nature Phys, 2(10):696–699, October 2006. Number: 10 Publisher: Nature Publishing Group.

[13] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, C. B. Schroeder, Cs. Tóth, J. Daniels, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, and E. Esarey. Multi-GeV Electron Beams from Capillary-Discharge-Guided Subpetawatt Laser Pulses in the Self-Trapping Regime. Phys. Rev. Lett., 113(24):245002, December 2014. Publisher: American Physical Society.

[14] F. Marteau, A. Ghaith, P. N'Gotta, C. Benabderrahmane, M. Valléau, C. Kitegi, A. Loulergue, J. Vétéran, M. Sebdaoui, T. André, G. Le Bec, J. Chavanne, C. Vallerand, D. Oumbarek, O. Cosson, F. Forest, P. Jivkov, J. L. Lancelot, and M. E. Couprie. Variable high gradient permanent magnet quadrupole (QUAPEVA). Applied Physics Letters, 111(25):253503, December 2017.

[15] 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(5):445–447, 1950.

[16] J. van Tilborg, S. Steinke, C. G. R. Geddes, N. H. Matlis, B. H. Shaw, A. J. Gonsalves, J. V. Huijts, K. Nakamura, J. Daniels, C. B. Schroeder, C. Benedetti, E. Esarey, S. S. Bulanov, N. A. Bobrova, P. V. Sasorov, and W. P. Leemans. Active Plasma Lensing for Relativistic Laser-Plasma-Accelerated Electron Beams. Phys. Rev. Lett., 115(18):184802, October 2015. Publisher: American Physical Society.

[17] R. Pompili, M. P. Anania, M. Bellaveglia, A. Biagioni, S. Bini, F. Bisesto, E. Brentegani, G. Castorina, E. Chiadroni, A. Cianchi, M. Croia, D. Di Giovenale, M. Ferrario, F. Filippi, A. Giribono, V. Lollo, A. Marocchino, M. Marongiu, A. Mostacci, G. Di Pirro, S. Romeo, A. R. Rossi, J. Scifo, V. Shpakov, C. Vaccarezza, F. Villa, and A. Zigler. Experimental characterization of active plasma lensing for electron beams. Appl. Phys. Lett., 110(10):104101, March 2017. Publisher: American Institute of Physics.

[18] J.-H. Röckemann, L. Schaper, S. K. Barber, N. A. Bobrova, G. Boyle, S. Bulanov, N. Delbos, K. Floettmann, G. Kube, W. Lauth, W. P. Leemans, V. Libov, A. R. Maier, M. Meisel, P. Messner, P. V. Sasorov, C. B. Schroeder, J. van Tilborg, S. Wesch, and J. Osterhoff. Direct measurement of focusing fields in active plasma lenses. Phys. Rev. Accel. Beams, 21(12):122801, December 2018. Publisher: American Physical Society.

[19] C. A. Lindstrøm, E. Adli, G. Boyle, R. Corsini, A. E. Dyson, W. Farabolini, S. M. Hooker, M. Meisel, J. Osterhoff, J.-H. Röckemann, L. Schaper, and K. N. Sjobak. Emittance Preservation in an Aberration-Free Active Plasma Lens. Phys. Rev. Lett., 121(19):194801, November 2018. Publisher: American Physical Society.

[20] H. Nakanishi, Y. Yoshida, T. Ueda, T. Kozawa, H. Shibata, K. Nakajima, T. Kurihara, N. Yugami, Y. Nishida, T. Kobayashi, A. Enomoto, T. Oogoe, H. Kobayashi, B. S. Newberger, S. Tagawa, K. Miya, and A. Ogata. Direct observation of plasma-lens effect. Phys. Rev. Lett., 66(14):1870–1873, April 1991. Publisher: American Physical Society.

[21] G. Hairapetian, P. Davis, C. E. Clayton, C. Joshi, S. C. Hartman, C. Pellegrini, and T. Katsouleas. Experimental demonstration of dynamic focusing of a relativistic electron bunch by an overdense plasma lens. Phys. Rev. Lett., 72(15):2403–2406, April 1994. Publisher: American Physical Society.

[22] J. S. T. Ng, P. Chen, W. Craddock, F. J. Decker, R. C. Field, M. J. Hogan, R. Iverson, F. King, R. E. Kirby, T. Kotseroglou, P. Raimondi, D. Walz, H. A. Baldis, P. Bolton, D. Cline, Y. Fukui, V. Kumar, C. Crawford, R. Noble, K. Nakajima, A. Ogata, and A. W. Weidemann. Results on Plasma Focusing of High Energy Density Electron and Positron Beams, August 2000. arXiv:physics/0008138.

[23] C. E. Clayton, B. E. Blue, E. S. Dodd, C. Joshi, K. A. Marsh, W. B. Mori, S. Wang, P. Catravas, S. Chattopadhyay, E. Esarey, W. P. Leemans, R. Assmann, F. J. Decker, M. J. Hogan, R. Iverson, P. Raimondi, R. H. Siemann, D. Walz, T. Katsouleas, S. Lee, and P. Muggli. Transverse Envelope Dynamics of a 28.5-GeV Electron Beam in a Long Plasma. Phys. Rev. Lett., 88(15):154801, April 2002. Publisher: American Physical Society.

[24] N. Barov and J. B. Rosenzweig. Propagation of short electron pulses in underdense plasmas. Phys. Rev. E, 49(5):4407–4416, May 1994.

[25] M. C. Thompson, H. Badakov, J. B. Rosenzweig, G. Travish, N. Barov, P. Piot, R. Fliller, G. M. Kazakevich, J. Santucci, J. Li, and R. Tikhoplav. Observations of low-aberration plasma lens focusing of relativistic electron beams at the underdense threshold. Physics of Plasmas, 17(7):073105, July 2010.

[26] R. Lehe, C. Thaury, E. Guillaume, A. Lifschitz, and V. Malka. Laser-plasma lens for laser-wakefield accelerators. Phys. Rev. ST Accel. Beams, 17(12):121301, December 2014.

[27] C. Thaury, E. Guillaume, A. Döpp, R. Lehe, A. Lifschitz, K. Ta Phuoc, J. Gautier, J.-P. Goddet, A. Tafzi, A. Flacco, F. Tissandier, S. Sebban, A. Rousse, and V. Malka. Demonstration of relativistic electron beam focusing by a laser-plasma lens. Nat Commun, 6(1):6860, April 2015. Number: 1 Publisher: Nature Publishing Group.

[28] Shyh-yuan Lee. Accelerator Physics (Fourth Edition). World Scientific Publishing Company, 2018.

[29] J. D. Jackson. Classical Electrodynamics. Wiley, Hoboken, 3rd edition edition, 1998.

[30] Christoph Bostedt, Sébastien Boutet, David M. Fritz, Zhirong Huang, Hae Ja Lee, Henrik T. Lemke, Aymeric Robert, William F. Schlotter, Joshua J. Turner, and Garth J. Williams. Linac Coherent Light Source: The first five years. Rev. Mod. Phys., 88(1):015007, March 2016. Publisher: American Physical Society.

[31] Richard L. Taylor. Magnetic Measurements of Quadrupole Focusing Magnets at SLAC. Technical Report SLAC-PUB-5621, Stanford Linear Accelerator Center, Stanford Univ., CA, 1991.

[32] E. D Courant and H. S Snyder. Theory of the alternating-gradient synchrotron. Annals of Physics, 3(1):1–48, January 1958.

[33] Robert Ariniello. Emittance Preservation in a Plasma Wakefield Accelerator. Ph.D. Thesis, University of Colorado at Boulder, United States – Colorado, 2022. ISBN: 9798802708729.

[34] J. Buon. Beam phase space and emittance. Technical report, France, 1992. LAL-RT–92-03 INIS Reference Number: 24057180.

[35] A. Sampath. Strong-field QED and collisional effects in electron beam-plasma interaction. Ph.D. Thesis, Ruperto-Carola-University of Heidelberg, 2020.

[36] Wolfgang Baumjohann and Rudolf A. Treumann. Basic Space Plasma Physics. World Scientific, 1997.

[37] Francis F. Chen. Introduction to Plasma Physics and Controlled Fusion. Springer, January 1984.

[38] Paul M. Bellan. Fundamentals of Plasma Physics. Cambridge University Press, July 2008.

[39] E. Esarey, C. B. Schroeder, and W. P. Leemans. Physics of laser-driven plasma-based electron accelerators. Rev. Mod. Phys., 81(3):1229–1285, August 2009. Publisher: American Physical Society.

[40] Rhon Keinigs and Michael E. Jones. Two-dimensional dynamics of the plasma wakefield accelerator. The Physics of Fluids, 30(1):252–263, January 1987.

[41] I. Blumenfeld, C. E. Clayton, F. J. Decker, M. J. Hogan, C. Huang, R. Ischebeck, R. H. Iverson, C. Joshi, T. Katsouleas, N. Kirby, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, E. Oz, R. H. Siemann, D. R. Walz, and M. Zhou. Scaling of the longitudinal electric field and transformer ratio in a nonlinear plasma wakefield accelerator. Phys. Rev. ST Accel. Beams, 13(11):111301, November 2010. Publisher: American Physical Society.

[42] Spencer J. Gessner. Demonstration of the Hollow Channel Plasma Wakefield Accelerator. Ph.D. Thesis, Stanford University, United States – California, 2016. ISBN: 9798662562363.

[43] Carl Andreas Lindstrøm. Emittance growth and preservation in a plasma-based linear collider. Ph.D. Thesis, University of Oslo, 2019.

[44] T Katsouleas, S Wilks, P Chen, J. M. Dawson, and J. J. Su. Beam Loading in Plasma Accelerators. Particle Accelerators, 22:81, 1987.

[45] James Albritton and Paul Koch. Cold plasma wavebreaking: Production of energetic electrons. The Physics of Fluids, 18(9):1136–1139, September 1975.

[46] W. Lu, C. Huang, M. Zhou, W. B. Mori, and T. Katsouleas. Nonlinear Theory for Relativistic Plasma Wakefields in the Blowout Regime. Phys. Rev. Lett., 96(16):165002, April 2006.

[47] W Lu, C Huang, M Zhou, M Tzoufras, F S Tsung, W B Mori, and T Katsouleas. A nonlinear theory for multidimensional relativistic plasma wave wakefields. Phys. Plasmas, 13(5):056709, 2006.

[48] A. A. Golovanov, I. Yu Kostyukov, L. Reichwein, J. Thomas, and A. Pukhov. Excitation of strongly nonlinear plasma wakefield by electron bunches. Plasma Phys. Control. Fusion, 63(8):085004, June 2021. Publisher: IOP Publishing.

[49] John M. Dawson. Nonlinear Electron Oscillations in a Cold Plasma. Phys. Rev., 113(2):383–387, January 1959. Publisher: American Physical Society.

[50] S. A. Yi, V. Khudik, C. Siemon, and G. Shvets. Analytic model of electromagnetic fields around a plasma bubble in the blow-out regime. Physics of Plasmas, 20(1):013108, January 2013.

[51] T. N. Dalichaouch, X. L. Xu, A. Tableman, F. Li, F. S. Tsung, and W. B. Mori. A multi-sheath model for highly nonlinear plasma wakefields. Physics of Plasmas, 28(6):063103, June 2021.

[52] Johannes Thomas, Igor Yu. Kostyukov, Jari Pronold, Anton Golovanov, and Alexander Pukhov. Non-linear theory of a cavitated plasma wake in a plasma channel for special applications and control. Physics of Plasmas, 23(5):053108, May 2016.

[53] A. A. Golovanov, I. Yu Kostyukov, A. M. Pukhov, and J. Thomas. Generalised model of a sheath of a plasma bubble excited by a short laser pulse or by a relativistic electron bunch in transversely inhomogeneous plasma. Quantum Electron., 46(4):295, April 2016. Publisher: IOP Publishing.

[54] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. Beam Loading in the Nonlinear Regime of Plasma-Based Acceleration. Phys. Rev. Lett., 101(14):145002, September 2008. Publisher: American Physical Society.

[55] T. Mehrling, J. Grebenyuk, F. S. Tsung, K. Floettmann, and J. Osterhoff. Transverse emittance growth in staged laser-wakefield acceleration. Phys. Rev. ST Accel. Beams, 15(11):111303, November 2012. Publisher: American Physical Society.

[56] M. Migliorati, A. Bacci, C. Benedetti, E. Chiadroni, M. Ferrario, A. Mostacci, L. Palumbo, A. R. Rossi, L. Serafini, and P. Antici. Intrinsic normalized emittance growth in laser-driven electron accelerators. Phys. Rev. ST Accel. Beams, 16(1):011302, January 2013. Publisher: American Physical Society.

[57] M. D. Litos, R. Ariniello, C. E. Doss, K. Hunt-Stone, and J. R. Cary. Beam emittance preservation using Gaussian density ramps in a beam-driven plasma wakefield accelerator. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 377(2151):20180181, August 2019. Publisher: Royal Society.

[58] R. Ariniello, C. E. Doss, K. Hunt-Stone, J. R. Cary, and M. D. Litos. Transverse beam dynamics in a plasma density ramp. Phys. Rev. Accel. Beams, 22(4):041304, April 2019. Publisher: American Physical Society.

[59] Chet Nieter and John R. Cary. VORPAL: a versatile plasma simulation code. Journal of Computational Physics, 196(2):448–473, May 2004.

[60] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby. A tensorial approach to computational continuum mechanics using object-oriented techniques. Computers in Physics, 12(6):620–631, November 1998. Publisher: American Institute of Physics.

[61] J. J. Su, T. Katsouleas, J. M. Dawson, and R. Fedele. Plasma lenses for focusing particle beams. Phys. Rev. A, 41(6):3321–3331, March 1990.

[62] A. Marocchino, M. P. Anania, M. Bellaveglia, A. Biagioni, S. Bini, F. Bisesto, E. Brentegani, E. Chiadroni, A. Cianchi, M. Croia, D. Di Giovenale, M. Ferrario, F. Filippi, A. Giribono, V. Lollo, M. Marongiu, A. Mostacci, G. Di Pirro, R. Pompili, S. Romeo, A. R. Rossi, J. Scifo, V. Shpakov, C. Vaccarezza, F. Villa, and A. Zigler. Experimental characterization of the effects induced by passive plasma lens on high brightness electron bunches. Appl. Phys. Lett., 111(18):184101, October 2017. Publisher: American Institute of Physics.

[63] R. Govil, W. P. Leemans, E. Yu. Backhaus, and J. S. Wurtele. Observation of Return Current Effects in a Passive Plasma Lens. Phys. Rev. Lett., 83(16):3202–3205, October 1999. Publisher: American Physical Society.

[64] C. A. Lindstrøm and E. Adli. Design of general apochromatic drift-quadrupole beam lines. Phys. Rev. Accel. Beams, 19(7):071002, July 2016. Publisher: American Physical Society.

[65] J. B. Rosenzweig and Pisin Chen. Beam optics of a self-focusing plasma lens. Phys. Rev. D, 39(7):2039–2045, April 1989. Publisher: American Physical Society.

[66] C. E. Doss, E. Adli, R. Ariniello, J. Cary, S. Corde, B. Hidding, M. J. Hogan, K. Hunt-Stone, C. Joshi, K. A. Marsh, J. B. Rosenzweig, N. Vafaei-Najafabadi, V. Yakimenko, and M. Litos. Laser-ionized, beam-driven, underdense, passive thin plasma lens. Phys. Rev. Accel. Beams, 22(11):111001, November 2019. Publisher: American Physical Society.

[67] C. Joshi, E. Adli, W. An, C. E. Clayton, S. Corde, S. Gessner, M. J. Hogan, M. Litos, W. Lu, K. A. Marsh, W. B. Mori, N. Vafaei-Najafabadi, B. O'shea, Xinlu Xu, G. White, and V. Yakimenko. Plasma wakefield acceleration experiments at FACET II. Plasma Phys. Control. Fusion, 60(3):034001, January 2018. Publisher: IOP Publishing.

[68] Katsunobu Oide. Synchrotron-Radiation Limit on the Focusing of Electron Beams. Phys. Rev. Lett., 61(15):1713–1715, October 1988. Publisher: American Physical Society.

[69] Kohji Hirata, Bruno Zotter, and Katsunobu Oide. Synchrotron-radiation limit of the luminosity in TeV linear colliders. Physics Letters B, 224(4):437–440, July 1989.

[70] P. Chen, K. Oide, A. M. Sessler, and S. S. Yu. Plasma-based adiabatic focuser. Phys. Rev. Lett., 64(11):1231–1234, March 1990. Publisher: American Physical Society.

[71] A. Aryshev, R. Ainsworth, T. Aumeyr, M. Bergamaschi, S. T. Boogert, P. Karataev, R. Kieffer, K. Kruchinin, T. Lefevre, S. Mazzoni, L. Nevay, N. Terunuma, and J. Urakawa. Submicron scale transverse electron beam size diagnostics methodology based on the analysis of optical transition radiation source distribution. J. Inst., 15(01):P01020, January 2020.

[72] Stewart T. Boogert, Grahame A. Blair, Gary Boorman, Alessio Bosco, Lawrence C. Deacon, Pavel Karataev, Alexander Aryshev, Masafumi Fukuda, Nobihiro Terunuma, Junji

Urakawa, Laura Corner, Nicolas Delerue, Brian Foster, David Howell, Myriam Newman, Rohan Senanayake, Roman Walczak, and Fred Ganaway. Micron-scale laser-wire scanner for the KEK Accelerator Test Facility extraction line. Phys. Rev. ST Accel. Beams, 13(12):122801, December 2010. Publisher: American Physical Society.

[73] Cosmin Ioan Blaga. Atoms and Molecules in Strong Midinfrared Laser Fields. Ph.D. Thesis, Stony Brook University, 2009.

[74] A M Perelomov, V S Popov, and M V Terent'Ev. Ionization of Atoms in an Aleternaing Electric Field. Soviet Physics JETP, 23(5):924, 1966.

[75] M. V. Ammosov, N. B. Delone, and V. P. Krainov. Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field. Sov. Phys. JETP, 64(1191):4, 1986.

[76] F. A. Ilkov, J. E. Decker, and S. L. Chin. Ionization of atoms in the tunnelling regime with experimental evidence using Hg atoms. J. Phys. B: At. Mol. Opt. Phys., 25(19):4005, October 1992.

[77] A. Couairon and A. Mysyrowicz. Femtosecond filamentation in transparent media. Physics Reports, 441(2):47–189, March 2007.

[78] David L. Bruhwiler, D. A. Dimitrov, John R. Cary, Eric Esarey, Wim Leemans, and Rodolfo E. Giacone. Particle-in-cell simulations of tunneling ionization effects in plasma-based accelerators. Physics of Plasmas, 10(5):2022–2030, May 2003. Publisher: American Institute of Physics.

[79] V S Popov. On the history of developing the theory of tunneling ionization in atoms and ions. Physics-Uspekhi, 42(7):733, 1999.

[80] Amnon Yariv. Quantum Electronics. Wiley, 3rd edition, 1989.

[81] Orazio Svelto. Principles of Lasers. Springer, 5th edition, 2010.

[82] Joseph W. Goodman. Introduction to Fourier Optics. Roberts and Company Publishers, 2005.

[83] J. A. Ratcliffe. Some Aspects of Diffraction Theory and their Application to the Ionosphere. Rep. Prog. Phys., 19(1):188, January 1956.

[84] A Behjat, G J Tallents, and D Neely. The characterization of a high-density gas jet. J. Phys. D: Appl. Phys., 30(20):2872–2879, October 1997.

[85] S. Semushin and V. Malka. High density gas jet nozzle design for laser target production. Review of Scientific Instruments, 72(7):2961–2965, July 2001. Publisher: American Institute of Physics.

[86] Krishna Pandey. CFD Analysis of Conical Nozzle for Mach 3 at Various Angles of Divergence with Fluent Software. International Journal of Chemical Engineering and Applications, Vol.1:179–185, August 2010.

[87] C. E. Doss, R. Ariniello, J. R. Cary, S. Corde, H. Ekerfelt, E. Gerstmayr, S. J. Gessner, M. Gilljohann, C. Hansel, B. Hidding, M. J. Hogan, A. Knetsch, V. Lee, K. Marsh, B. O'Shea, P. San Miguel Claveria, D. Storey, A. Sutherland, C. Zhang, and M. D. Litos. Underdense plasma lens with a transverse density gradient. Phys. Rev. Accel. Beams, 26(3):031302, March 2023. Publisher: American Physical Society.

[88] S. Kuschel, D. Hollatz, T. Heinemann, O. Karger, M. B. Schwab, D. Ullmann, A. Knetsch, A. Seidel, C. Rödel, M. Yeung, M. Leier, A. Blinne, H. Ding, T. Kurz, D. J. Corvan, A. Sävert, S. Karsch, M. C. Kaluza, B. Hidding, and M. Zepf. Demonstration of passive plasma lensing of a laser wakefield accelerated electron bunch. Phys. Rev. Accel. Beams, 19(7):071301, July 2016. Publisher: American Physical Society.

[89] G. Stupakov. Generation of a wakefield undulator in plasma with transverse density gradient. Physics of Plasmas, 24(11):113110, November 2017.

[90] Chaojie Zhang, Zan Nie, Yipeng Wu, Mitchell Sinclair, Chen-Kang Huang, Ken A. Marsh, and Chan Joshi. Ionization induced plasma grating and its applications in strong-field ionization measurements. Plasma Phys. Control. Fusion, 63(9):095011, August 2021. Publisher: IOP Publishing.

[91] W. K. H. Panofsky and W. A. Wenzel. Some Considerations Concerning the Transverse Deflection of Charged Particles in Radio-Frequency Fields. Review of Scientific Instruments, 27(11):967–967, November 1956.

[92] S. Vaganian and H. Henke. The Panofsky-Wenzel Theorem and General Relations for the Wake Potential. Particle Accelerators, 48:239–242, 1995.

[93] E. Adli, S. Corde, R. J. England, J. Frederico, S. J. Gessner, S. Li, Z. Wu, M. J. Hogan, M. D. Litos, W. An, C. E. Clayton, K. A. Marsh, W. Lu, W. Mori, C. Joshi, N. Vafaei-Najafabadi, and P. Muggli. First Results from the Electron Hose Instability Studies in FACET. Proc. IPAC'12, New Orleans, Louisiana, USA, pages 43–45, 2012.

[94] B. Miao, J. E. Shrock, L. Feder, R. C. Hollinger, J. Morrison, R. Nedbailo, A. Picksley, H. Song, S. Wang, J. J. Rocca, and H. M. Milchberg. Multi-GeV Electron Bunches from an All-Optical Laser Wakefield Accelerator. Phys. Rev. X, 12(3):031038, September 2022.

[95] J. E. Shrock, B. Miao, L. Feder, and H. M. Milchberg. Meter-scale plasma waveguides for multi-GeV laser wakefield acceleration. Physics of Plasmas, 29(7):073101, July 2022.

[96] Y. Ma, D. Seipt, S. J. D. Dann, M. J. V. Streeter, C. A. J. Palmer, L. Willingale, and A. G. R. Thomas. Angular streaking of betatron X-rays in a transverse density gradient laser-wakefield accelerator. Physics of Plasmas, 25(11):113105, November 2018.

[97] P. Muggli, S. Lee, T. Katsouleas, R. Assmann, F. J. Decker, M. J. Hogan, R. Iverson, P. Raimondi, R. H. Siemann, D. Walz, B. Blue, C. E. Clayton, E. Dodd, R. A. Fonseca, R. Hemker, C. Joshi, K. A. Marsh, W. B. Mori, and S. Wang. Collective refraction of a beam of electrons at a plasma-gas interface. Phys. Rev. ST Accel. Beams, 4(9):091301, September 2001.

[98] G. Wittig, O. Karger, A. Knetsch, Y. Xi, A. Deng, J. B. Rosenzweig, D. L. Bruhwiler, J. Smith, G. G. Manahan, Z.-M. Sheng, D. A. Jaroszynski, and B. Hidding. Optical plasma torch electron bunch generation in plasma wakefield accelerators. Phys. Rev. ST Accel. Beams, 18(8):081304, August 2015.

[99] Carl A. Lindstrøm. Staging of plasma-wakefield accelerators. Phys. Rev. Accel. Beams, 24(1):014801, January 2021.

[100] C. Emma, A. Edelen, M. J. Hogan, B. O'Shea, G. White, and V. Yakimenko. Machine learning-based longitudinal phase space prediction of particle accelerators. Phys. Rev. Accel. Beams, 21(11):112802, November 2018. Publisher: American Physical Society.

[101] S. Z. Green, E. Adli, C. I. Clarke, S. Corde, S. A. Edstrom, A. S. Fisher, J. Frederico, J. C. Frisch, S. Gessner, S. Gilevich, P. Hering, M. J. Hogan, R. K. Jobe, M. Litos, J. E. May, D. R. Walz, V. Yakimenko, C. E. Clayton, C. Joshi, K. A. Marsh, N. Vafaei-Najafabadi, and P. Muggli. Laser ionized preformed plasma at FACET. Plasma Phys. Control. Fusion, 56(8):084011, July 2014. Publisher: IOP Publishing.

[102] V. Yakimenko, L. Alsberg, E. Bong, G. Bouchard, C. Clarke, C. Emma, S. Green, C. Hast, M. J. Hogan, J. Seabury, N. Lipkowitz, B. O'Shea, D. Storey, G. White, and G. Yocky. FACET-II facility for advanced accelerator experimental tests. Phys. Rev. Accel. Beams, 22(10):101301, October 2019. Publisher: American Physical Society.

[103] Roswell Lee and Martin Lampe. Electromagnetic Instabilities, Filamentation, and Focusing of Relativistic Electron Beams. Phys. Rev. Lett., 31(23):1390–1393, December 1973. Publisher: American Physical Society.

[104] Simon Hooker and Colin Webb. Laser Physics. Oxford University Press, August 2010.

[105] V.L. Ginzburg and I.M. Frank. Radiation of a uniformly moving electron due to its transition from one medium into another. Journal of Physics (USSR), 9:353–362, 1945.

[106] Selina Green, Mark Hogan, Nate Lipkowitz, Brendan O'Shea, Glen White, Vitaly Yakimenko, and Gerald Yocky. Beam Diagnostic Challenges for FACET-II. In Proceedings of the 6th Int. Beam Instrumentation Conf., Grand Rapids, MI, 2017.

[107] N. Kirby, M. Berry, I. Blumenfeld, M. J. Hogan, R. Ischebeck, and R. Siemann. Emittance growth from Multiple Coulomb Scattering in a plasma wakefield accelerator. In 2007 IEEE Particle Accelerator Conference (PAC), pages 3097–3099, June 2007. ISSN: 2152-9582.