Title

Controlling Electron Acceleration in Underdense and Overdense Laser–Plasma Interactions to Generate X-rays for Probing High Energy Density Material

Author

Miller, Kyle G

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Controlling Electron Acceleration in Underdense and Overdense Laser–Plasma Interactions to Generate X-rays for Probing High Energy Density Material

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Physics

by

Kyle G. Miller

2022

ABSTRACT OF THE DISSERTATION

Controlling Electron Acceleration in Underdense
and Overdense Laser–Plasma Interactions
to Generate X-rays for Probing
High Energy Density Material

by

Kyle G. Miller
Doctor of Philosophy in Physics
University of California, Los Angeles, 2022
Professor Warren B. Mori, Chair

There is interest in using short-pulse x-rays that are small in source size, broad in energy spectrum, and high in photon number to probe and visualize the evolution of hot, dense material for both research and industrial applications. One method to produce such x-rays is to collide an energetic electron beam with a high-Z material, which will then emit bremsstrahlung radiation with many of the desired source characteristics. In this dissertation we study the physical processes of generating energetic electrons from laser–plasma interactions in both underdense and overdense plasmas.

These laser–plasma interactions are nonlinear and kinetic in nature. Therefore, the particle-in-cell (PIC) algorithm is often the tool of choice for the simulations discussed within this dissertation, with length and time scales on the order of a millimeter and picosecond, respectively. Such simulations require the use of massively parallel computers. However, these simulations often suffer from having a large concentration of particles processed by relatively

few computing elements, leading to decreases in performance due to a computational load imbalance. We present a dynamic load balancing technique for the PIC algorithm that effectively balances computational load across distributed-memory processes, in addition using a hybrid shared-memory scheme to increase scalability with shared-memory thread number by an order of magnitude and boost overall performance by a factor of two. Another useful PIC algorithm relevant to this work invokes a cylindrical geometry and azimuthal mode decomposition to yield proper three-dimensional geometric effects at the computational cost of a two-dimensional simulation. We also discuss improvements to this algorithm, where modifications to the particle initialization and field solver at the cylindrical axis eliminate spurious electromagnetic fields at the axis that have long been observed for this method.

The second part of this dissertation explores the mechanism of direct laser acceleration (DLA) in laser-based plasma acceleration. This process occurs when the channel-guided laser fields overlap electrons either in the plasma wave wake or within an ion channel, and the frequency of the electron transverse motion matches the Doppler-shifted laser frequency. We first utilize the cylindrical mode decomposition to more accurately account for the energy gain from the DLA process compared to traditional methods, then show that laser wakefield accelerators (LWFAs) in both the self-modulated (SM-LWFA) and bubble regimes exhibit comparable contributions in energy from the wakefields and DLA process for the most energetic electrons. A customized finite-difference Maxwell field solver is then presented that corrects the dispersion relation of light in vacuum and removes a time-discretization error in the Lorentz force compared to the standard PIC algorithm. This solver is especially valuable when investigating DLA, and simulations using the customized solver demonstrate better agreement with experiment and with numerically integrated equations of motion. Single-particle motion is analyzed to study resonant motion in the DLA process, where electrons are observed to gain significant energy from laser fields but do not readily transition between different orders of resonance to gain further energy. We also simulate the motion of an electron probe beam propagating across an LWFA perpendicular to the laser propagation

direction, which is timed with the laser pulse and measured far from the plasma to image the dynamics of the plasma wave wake. Although some qualitative agreement is observed between simulation and experiment, further investigation is needed to discern wakefield properties from the radiograph image alone.

In the last part of this dissertation, we present simulations of laser–solid interactions to investigate the dynamics of energetic electron generation in the density upramp before an overdense plasma. These electrons propagate through the target and are then collided with a high-Z material to emit bremsstrahlung radiation. We first detail the requisite simulation techniques to correctly model this process, namely the splitting of energetic macro-particles to reduce enhanced wakefields, an extended particle absorber to prevent reflux at the boundary and a large transverse domain size to resolve long-wavelength magnetic field modes in the low-density plasma. A series of simulations are then carried out with varied laser amplitude and duration at constant energy to determine the laser configuration that yields the largest dose of few-MeV x-rays. We find that the high-amplitude laser pulses generate higher-temperature electron spectra, which in turn produce more x-rays at the desired energy. In addition, the shortest pulses generate many energetic electrons before the formation of self-generated magnetic fields, resulting in more directional beams of electrons and x-rays.

The dissertation of Kyle G. Miller is approved.

Troy A. Carter

George J. Morales

Chandra J. Joshi

Warren B. Mori, Committee Chair

University of California, Los Angeles

2022

To my wife and sons,
for whom I am doing all of this
and who supported me at work
and in the home office for many years.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank my advisor, Warren Mori, for helping me to further my education throughout these last six years. He made a consistent effort to help me pursue the topics that most interested me, and he made available many projects and collaborations that have formed the foundation for my scientific career. I am also grateful to the other scientists in our group—Frank Tsung, Viktor Decyk, Benjamin Winjum and Paulo Alves—that made themselves available for many a helpful discussion on an array of topics. Special thanks as well to Ricardo Fonseca for his careful watch over much of my software development throughout the years, for helping to educate me about the nuances of parallel computing and to establish habits that will serve me for years to come.

Special thanks to my many collaborators at various institutions who have aided me in producing quality work, expanded my scientific knowledge and performed excellent experiments for me to work with. Thanks to Félicie Albert, Nuno Lemos, Scott Wilks, Andreas Kemp, Dean Rusby and Paul King from Lawrence Livermore National Laboratory, Navid Vafaei-Najafabadi from Stony Brook University, Chan Joshi from the University of California, Los Angeles and Jessica Shaw from the Laboratory for Laser Energetics at the University of Rochester for the many hours spent and quality work performed to help me succeed in my research.

I am also thankful to the many students and postdocs that I have worked with over the years. They have given me new perspectives and insight into my work and professional career. I am thankful to Paul Elias, Archis Joglekar, Han Wen, Xinlu Xu and Fei Li for their programming wisdom and examples to me as computational physicists as I was first learning the ropes. I would like to especially thank Adam Tableman for his apparent extreme devotion to the rising generation of physicists in our group and for many (often very opinionated) discussions about parallel computing that have helped me form a solid base of knowledge and experience. It was an absolute pleasure to work along with the many students I have

interacted with along the way, including Asher Davidson, Joshua May, Anton Helm, Thamine Dalichaouch, Lance Hildebrand, Yujian Zhao, Roman Lee, Jacob Pierce, Sarah Chase and Qianqian Su. I would like to especially thank Roman Lee for the many long hours sustained together, staring at a computer screen and working on the code base that would yield the bulk of a chapter in this dissertation.

I cannot overstate the importance and impact that my wife, Hannah Miller, has played in my life. She is the reason that I wake up in the morning. She is my perpetual cheerleader, my constant in life, and the one person in this world who seems to understand me better than I understand myself. I am forever grateful for the joyous countenances of our two sons that grace our home, and I am glad that they have survived graduate school with me. To many more happy years.

VITA

2009–2015 B.S. in physics, *summa cum laude*, Brigham Young University, Provo, UT.

2015–2016 M.S. in physics, emphasis in acoustics, Brigham Young University.

2016–2022 Graduate Student Researcher, University of California, Los Angeles, Los Angeles, CA.

PUBLICATIONS

X. Xu, F. Li, F. S. Tsung, K. G. Miller, V. Yakimenko, M. J. Hogan, C. Joshi and W. B. Mori, "Generation of ultrahigh-brightness, pre-bunched beams from a plasma cathode for Xray free-electron lasers," *Nat. Commun.*, (accepted 2022).

K. G. Miller, J. May, F. Fiuza, and W. B. Mori, "Extended particle absorber for efficient modeling of intense laser–solid interactions," *Phys. Plasmas*, vol. 28, p. 112702 (2021).

F. Li, V. K. Decyk, K. G. Miller, A. Tableman, F. S. Tsung, M. Vranic, R. A. Fonseca and W. B. Mori, "Accurately simulating nine-dimensional phase space of relativistic particles in strong fields," *J. Comput. Phys.*, vol. 438, p. 110367 (2021).

J. L. Shaw, M. A. Romo-Gonzalez, N. Lemos, P. M. King, G. Bruhaug, K. G. Miller, C. Dorrer, B. Kruschwitz, L. Waxer, G. J. Williams, M. V. Ambat, M. M. McKie, M. D. Sinclair, W. B. Mori, C. Joshi, H. Chen, J. P. Palastro, F. Albert and D. H. Froula, "Microcoulomb ($0.7 \pm \frac{0.4}{0.2}$ μC) laser plasma accelerator on OMEGA EP," *Sci. Rep.*, vol. 11, p. 7498 (2021).

P. M. King, K. Miller, N. Lemos, J. L. Shaw, B. F. Kraus, M. Thibodeau, B. M. Hegelich, J. Hinojosa, P. Michel, C. Joshi, K. A. Marsh, W. Mori, A. Pak, A. G. R. Thomas and F. Albert, "Predominant contribution of direct laser acceleration to high-energy electron spectra in a low-density self-modulated laser wakefield accelerator," *Phys. Rev. Accel. Beams*, vol. 24, p. 011302 (2021).

K. G. Miller, R. P. Lee, A. Tableman, A. Helm, R. A. Fonseca, V. K. Decyk and W. B. Mori, "Dynamic load balancing with enhanced shared-memory parallelism for particle-in-cell codes," *Comput. Phys. Commun.*, vol. 259, p. 107633 (2021).

F. Li, K. G. Miller, X. Xu, F. S. Tsung, V. K. Decyk, W. An, R. A. Fonseca and W. B. Mori, "A new field solver for modeling of relativistic particle-laser interactions using the particle-in-cell algorithm," *Comput. Phys. Commun.*, vol. 258, p. 107580 (2021).

K. G. Miller and K. L. Gee, "Model-scale jet noise analysis with a single-point, frequency-domain nonlinearity indicator," *J. Acoust. Soc. Am.*, vol. 143, pp. 3479-3492 (2018).

K. G. Miller, K. L. Gee, and B. O. Reichman, "Asymptotic behavior of a frequency-domain nonlinearity indicator for solutions to the generalized Burgers equation," *J. Acoust. Soc. Am.*, vol. 140, pp. EL522-EL527 (2016).

K. G. Miller, M. Meehan, R. L. Spencer, and J. S. Colton, "Resonance of Complex Cylindrically Symmetric Cavities Using an Eigenfunction Expansion in Empty Cavity Modes," *IEEE Trans. Microw. Theory Techn.*, vol. 64, pp. 3113-3120 (2016).

CHAPTER 1

Introduction

As laser intensities available in the laboratory continue to increase year after year [2, 3, 4], so also does the study of materials at high energy density (HED), also referred to as high energy density science (HEDS). In the field of plasma physics, this research ranges from the study of relativistic shocks [5, 6, 7] to ultra-relativistic particle beams produced from high-intensity lasers [8, 9]. However, extreme conditions and short lifetimes of these materials make them very difficult to image and diagnose. For this purpose, it is desirable to perform non-destructive material imaging through x-ray radiography. The x-ray sources for these applications should ideally be short pulsed (fs–ns), have small source size (1–100s of μm), exhibit a broad energy spectrum (50 keV to few MeV) and yield a large number of photons ($> 10^{10}$). These light sources can then be used to create "movies" of dense, rapidly moving objects either for research purposes [10] or within industry [11, 12, 13].

One efficient method for generating such x-rays is to collide an energetic (1–100 MeV) electron beam with a high-Z material to produce bremsstrahlung radiation [10, 14]. However, conventional electron sources are not readily available in many laboratories, and the temporal and spatial resolutions of the produced electron bunches are not always well suited for HEDS diagnostic purposes. Instead, various methods have been proposed to utilize the powerful laser systems available at many facilities to produce the energetic electron beams that can then emit the desired x-rays [15, 16]. A high-intensity laser incident on a solid target can accelerate electrons in the density upramp region, which then emit bremsstrahlung x-rays as they propagate either through the remainder of the solid target [10, 17] or through a

high-Z material located behind the thin target [14, 18]. Energetic electrons created in an underdense plasma from a laser wakefield accelerator (LWFA) [19, 20, 11] or self-modulated laser wakefield accelerator (SM-LWFA) [21] can also be used for the x-ray emission. In each of these cases, the high-intensity laser fields result in laser–plasma interactions that are highly nonlinear, necessitating fully kinetic simulations to predict and understand the inherent physical processes. The scope of this dissertation is to both explore and learn to control the electron acceleration mechanisms present in relativistic laser–plasma interactions via particle-in-cell simulation with the goal of producing x-ray sources relevant to HEDS applications.

1.1 Methodology: particle-in-cell simulation

The particle-in-cell (PIC) algorithm [22, 23] aims to simulate the interaction of many charged particles with their self-consistent electromagnetic fields. Groups of charged particles are represented by macro-particles that have the same charge-to-mass ratio as their individual counterparts, but are of finite size. (We note that it is possible for a macro-particle to represent a single, but finite-size particle.) These macro-particles carry a charge and "continuous" (to floating-point precision) positions and momenta. Electric and magnetic fields, along with the current and charge are defined on a regular grid, with the various components in most cases offset from one another in time and space according to what is now often referred to as the Yee mesh [24].

A simulation is initialized with a distribution of particles (each with some initial charge, position and momentum) and some set of electric and magnetic fields. The forces on individual macro-particles are calculated from the discretized fields by interpolating some number of nearby field quantities onto each particle. The manner in which this is done depends on the order of the shape function, which is generally calculated with B-splines and can be of order linear, quadratic, etc. Once these forces are determined, the continuous particle momenta

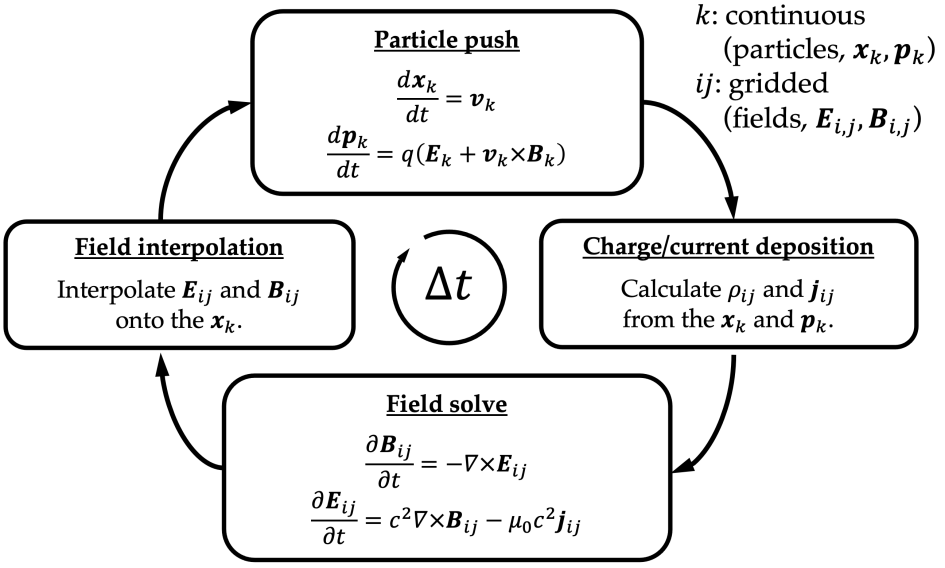

Figure 1.1: Loop depicting a single time step in the particle-in-cell (PIC) algorithm.

are updated according to the Lorentz force equation, most often via the Boris push [25]. Particle positions are then updated assuming that particles drift in a straight line for a single time step. Next, the particle shape functions[1] are used to deposit current into nearby grid quantities, after which the electric and magnetic fields are solved based on Faraday's and Ampere's laws.[2] The field solve is usually performed using a finite-difference scheme or in Fourier space using fast Fourier transforms. This completes the PIC loop, depicted in Fig. 1.1, which is executed many times over the course of a simulation.

There are many nuances to this algorithm, including methods for computing particle momenta [32, 33, 34], solving for electromagnetic fields [35, 36, 37], handling field and particle

[1]The particle "shapes" used to deposit particle data to the grid (e.g., current and charge deposit) and interpolate grid data (e.g., fields) to the particles can be the same or different. When they are the same, the algorithm can be momentum conserving.

[2]If the electric and magnetic fields have the correct divergences at initialization, they will retain this property throughout the entire simulation as long as the continuity equation is satisfied by the charge and current. Some charge deposition schemes do not obey the continuity equation, and the electric field must in turn be corrected by one of various schemes [26, 23, 27]. OSIRIS, however, deposits the current in a way that satisfies the continuity equation, called a charge-conserving current deposit [28, 29, 30, 31], obviating the need to solve Gauss's law.

boundary conditions [38, 23, 39], and parallelization schemes [40, 41, 42]. In this dissertation we use the PIC code OSIRIS [28, 29], which has the ability to use high-order particle shapes (up to quartic), many different implementations for updating particle momenta, many finite-difference field solvers [36, 37], and also scales well to over one million cores [28, 43]. The OSIRIS development community is very active, and many new features are added each year.

In addition, there is much activity aimed at improving the performance of the OSIRIS software. Many pertinent physical problems simulated with PIC codes involve concentrated bunches of particles, e.g., laser wakefield acceleration (discussed later in detail) or magnetic reconnection with pair production [44]. Performance can suffer when relatively few computing resources on a parallel computer are responsible for processing a majority of particle trajectories. Shared-memory parallelization has been utilized to balance computational load across cores and decrease simulation runtime [43]. Other PIC models can also greatly increase performance by making simplifying assumptions regarding problem constraints. For example, running in a cylindrically symmetric configuration [30] or in a Lorentz-boosted frame [45, 46] can yield similar results to standard methods at a great reduction in computational cost; a combination of these two algorithms can in some cases achieve speedups on the order of 100,000 [47]. Some of these new features are pertinent to this research and will be described in later chapters.

With today's computing technology, the PIC algorithm is often the tool of choice for simulating laser–plasma interactions for spatial scales on the order of millimeters and tens of picoseconds in duration, and some simulations are even now performed to investigate multi-speckled laser beams [48]. In particular, the physics problems discussed in this dissertation are all amenable to PIC simulation. Herein we will discuss many improvements to the PIC algorithm, as well as show results from many large-scale simulations.

1.2 Electron acceleration driven by lasers in underdense plasmas

Electromagnetic radiation will only propagate through a plasma if it is oscillating at a frequency greater than the electron plasma frequency, $\omega_p = \sqrt{n_e e^2/\varepsilon_0 m_e}$, where n_e, e and m_e are the density, charge and rest mass of an electron, respectively, and ε_0 is vacuum permittivity. For a laser of frequency ω_0, we define the critical density, n_c, to be the electron density where $\omega_p = \omega_0$, i.e., $n_c \equiv \omega_0^2 \varepsilon_0 m_e/e^2$. Below this density laser light will propagate through the plasma (underdense), whereas it will not for larger densities (overdense). At high laser intensities, electrons within the plasma oscillate at energies that can be highly relativistic. In these cases, there are relativistic mass corrections for individual electrons and fluid quantities. These corrections can reduce the plasma frequency, since it is proportional to the proper density, n_e/γ, where γ is the relativistic factor for an electron. It can be shown [49] that γ scales roughly as $\sqrt{1 + \frac{1}{2}(eA_0/m_e c)^2}$, where A_0 is the peak vector potential of a linearly polarized laser. The normalized vector potential (used later) is defined as $a_0 \equiv eA_0/m_e c$, which is also often called v_{osc}/c. Thus, a more general expression for the critical density is $n_c = \gamma \omega_0^2 \varepsilon_0 m_e/e^2$.

A laser pulse propagating through an underdense plasma can push electrons away from the laser axis through the ponderomotive force, after which they are attracted back to the axis by the focusing force from the exposed (and relatively stationary) ions. Depending on the laser pulse amplitude and duration, the expelled electrons can come back to the axis at nearly the same time, forming what is called a plasma wave wake. The extreme concentration of this negative charge results in a large electric field, which can be supported by the plasma [50, 51] and sustained over long distances in what is known as a laser wakefield accelerator (LWFA) [52]. The accelerating fields sustainable in a plasma of density $10^{18}\,\mathrm{cm}^{-3}$ are roughly $100\,\mathrm{GV/m}$, several orders of magnitude larger than the breakdown limit of conventional accelerators [53, 54].

Many alternatives to traditional LWFA have been proposed for plasma-based accelerators,

including the plasma beat wave accelerator (PBWA) [55, 56]—where two laser pulses are used with frequencies whose difference matches the electron plasma frequency—the self-modulated laser wakefield accelerator (SM-LWFA) [55, 57, 58, 59]—where a long-duration laser pulse breaks into a train of short pulses on the order of the plasma wavelength—or the plasma wakefield accelerator (PWFA) [60, 59, 61]—where an energetic particle driver is used to excite the plasma wake. In addition, the physical response of the plasma differs greatly depending on the normalized vector potential of a laser driver, $a_0 \simeq 8.6 \times 10^{-10}\sqrt{I_0\,[\mathrm{W/cm^2}]}\lambda\,[\mu\mathrm{m}]$. For a low-amplitude laser with $a_0 \ll 1$, the plasma wave behaves linearly and the response can be mostly described analytically [62, 63]. However, with increased laser amplitude ($a_0 \gtrsim 1$) the problem becomes highly nonlinear [49] and becomes very difficult to analytically model in three dimensions. (For this reason PIC simulations are especially valuable for investigating the nonlinear optics of and plasma response to high-intensity lasers.) If the laser amplitude is large enough, it can completely expel all electrons near the laser and form a large spherical cavity [64, 65, 66]; this is commonly referred to as the blowout or bubble regime. An analytic theory of the bubble radius and dephasing length has been developed for this regime [66, 67], providing valuable guidance for constructing and simulating such an LWFA.

For some of the cases mentioned, a portion of the laser may overlap particles being accelerated by the wakefields. This commonly occurs either for a short-pulse laser that completely fills the first plasma bucket [65, 68, 69, 1] or for a long-pulse laser that overlaps many bunches of electrons, as in SM-LWFA [70, 71, 72, 73, 74]. When such overlap is present, there can be cooperative acceleration mechanisms between the wakefield and laser field. Electrons near an exposed ion column experience a focusing force that causes them to execute (relativistic) harmonic motion at the betatron frequency [75], $\omega_\beta = \omega_p/\sqrt{2\gamma}$. The transverse oscillation is accompanied by emission of betatron radiation [75, 76, 77], which is emitted in the forward direction and accumulates from many electrons. If the Doppler-shifted laser frequency seen by an electron matches the betatron frequency, a resonance can occur where the electron gains energy from the transverse laser electric field that is rotated to

forward momentum via the $\vec{v} \times \vec{B}$ force. This process is known as direct laser acceleration (DLA) [78, 70, 79, 1]. The energy gain from the DLA process can be similar in magnitude to the energy gained from the accelerating wakefields, and distinguishing the relative contribution from of the differing processes is an important research question both in experiment and simulation [68, 80, 1, 74]. In addition, for a laser that is sufficiently long in duration, the ions can also be pushed away from the laser axis, leaving the electrons to oscillate in and gain energy from the laser fields that are guided in the plasma channel [81, 73].

Although experimental work has demonstrated the success of many of these plasma-based acceleration approaches, much work remains to optimize and refine the quality of accelerated electron beams [82, 9] to be used as x-ray light sources [83, 84, 85]. In particular, being able to better understand and sustain betatron resonance could increase both overall electron energy as well as the emitted betatron radiation. Simulations [85, 86, 87, 88] and some recent experiments [89, 90] have shown that plasma-based acceleration can produce the required beam brightness for building a compact free-electron x-ray laser [9]. However, plasma-based accelerators can also provide compact sources of electrons to generate ultra-compact, incoherent but directional x-ray sources [85, 91, 92].

1.3 Electron acceleration from laser–solid interactions

For a plasma with density larger than the critical density, the plasma behaves as a solid to the laser, and light will be reflected at the critical-density interface (or absorbed into the plasma) [93, 94]. One method for generating energetic electrons is to illuminate a solid target with a high-intensity laser. For a pre-ionized target with a low-density plasma ramp in front of the target, the incident and reflected light waves interacting with the pre-plasma can stochastically accelerate electrons to very high energies and direct them forward into the target. This process is complicated, as the surface deforms and various acceleration mechanisms are active [95, 96, 97, 98]. The system dynamics also vary greatly based on

the physical size of the target. For example, targets that are thin in the laser propagation direction can generate strong electric fields due to exiting electrons at the rear of the target, fields that are then responsible for accelerating heavy ions forward at high energies [99, 100]. In some applications, energetic electrons that propagate through the target can be collided with a high-Z material to produce bremsstrahlung x-rays [10, 14]. Much work has been devoted to increasing the conversion efficiency of this process, including the use of cone targets [101, 102], introducing an angle of incidence [103, 94] and using multiple pulses slightly delayed from one another [104].

The exact mechanisms for the stochastic electron acceleration in these experiments are still not well understood [14, 105]. Perhaps a more open research question is how to correctly predict the electron temperature produced from such a laser–solid interaction. Many theories and parameter fits have been calculated for various experiments, but results appear to vary greatly based on laser amplitude and duration, as well as the plasma configuration [105]. In addition, simulation of these processes is difficult, and simplifying assumptions often have to be made. However, care must be taken to not exclude crucial physics processes with such assumptions; for example, one-dimensional simulations have been shown to completely miss one acceleration process at a steep plasma interface [97]. There is clearly much research that remains to be done in this area both to more reliably simulate the laser–solid interactions as well as to better understand the physical processes involved in electron acceleration.

1.4 Dissertation outline

The dissertation is organized as follows. In Chapter 2 we present computational developments in the PIC code Osiris. The first divides the simulation space into small regions ("tiles") to better balance computational load, and an order-of-magnitude improvement in scalability with thread count is observed due to a new method of shared-memory parallelization. The second set of developments improves the charge initialization and field solve near the axis of

cylindrically symmetric OSIRIS simulations, and also gives an implementation of an open boundary condition for electromagnetic fields at left and right boundaries of the cylindrical geometry. Chapter 3 discusses betatron resonance and the impact of direct laser acceleration in laser wakefield accelerators. A more accurate method of computing the energy gain from direct laser acceleration is given, after which full-scale simulations show the interplay between the accelerating wakefields and laser fields on high-energy electrons. A customized field solver is also developed, then used to give better agreement with experiment for a laser wakefield accelerator simulation. Electron radiography is then studied for use with a laser wakefield accelerator. In Chapter 4 we present challenges that were overcome to properly simulate large-scale laser–solid interactions, after which we present simulations with laser amplitude and duration varied by more than three orders of magnitude. The electron and x-ray spectra are extracted and analyzed, showing a clear trend in desired x-ray generation. Finally, concluding remarks are made in Chapter 5, along with direction for future work.

CHAPTER 2

Software and algorithm development for high-fidelity particle-in-cell simulations

To model full-scale, long-time simulations relevant to this dissertation—such as direct laser acceleration in self-modulated laser wakefield accelerators or hot-electron generation from picosecond laser–plasma interactions in overdense plasmas—we require software modifications to the OSIRIS framework that will ensure accuracy and feasibility. Large groups of macro-particles often bunch together in simulations of plasma-based acceleration, which can significantly decrease performance due to an imbalance in computational load (i.e., few computing elements processing a majority of particles). When simulating laser–plasma interactions for picosecond lasers ($\sim$ 1000 laser periods/wavelengths) propagating in underdense plasmas, the required domain size and number of particles can result in prohibitively demanding simulations in three dimensions. The quasi-3D OSIRIS algorithm [30] can provide results with the correct three-dimensional geometric effects at the cost of a two-dimensional simulation, though care must be taken to ensure that noise or spurious field growth does not occur near the cylindrical axis over time. In this chapter, we describe first a new approach to dynamic load balancing that combines two different shared-memory parallelization models, giving an order-of-magnitude increase in scalability with thread number compared to previous implementations. We then present several modifications to the quasi-3D algorithm that correct long-standing issues in the charge, current and fields near the cylindrical axis.

2.1 Dynamic load balancing in OSIRIS via tiles

2.1.1 Introduction

The particle-in-cell (PIC) algorithm [22, 23] is widely used to study interesting problems where discrete particles or agents interact through fields. The PIC algorithm has thus been widely used in the kinetic modeling of plasmas, where the fields can either be electrostatic or electromagnetic. However, the nature of tracking individual particles over long periods of time makes the PIC algorithm computationally expensive, requiring the use of large-scale high-performance computing (HPC) resources. With the advent of exascale computing, HPC architectures are undergoing rapid change; since 2004, clock rates have stabilized and growth on top-ranked systems has come almost entirely from increased parallelism (often within a CPU). The PIC algorithm is sensitive to load imbalance on excessively parallel machines because simulation particles move significant distances and may accumulate on a fraction of the computing resources. As the scale of massively parallel computing architectures continues to increase, the final push toward exascale and beyond will require significant adaptation of software to take advantage of the increased parallelism available in the hardware.

Distributing computational load evenly across resources can be achieved through multiple levels of parallelism. At the highest level one can parallelize across distributed-memory processing elements (PEs). By this definition a processing element and a node are not the same thing, as one can use various PEs within a single node. Parallelism on this level is often implemented via the Message Passing Interface (MPI) standard, and for clarity, in this section PE always refers to an MPI process. In addition, the increasing number of shared-memory CPUs inside compute nodes, as well as the increasing number of cores inside today's CPUs, allows for parallelism within a processing element (e.g., via OpenMP or Pthreads). Finally, the use of many-core accelerators from multiple vendors such as Graphical Processing Units (GPUs) also allows for parallelism via CUDA, OpenACC, HIP, or SYCL/DPC++.

In this section we present developments to improve the parallel scalability of the particle-

in-cell (PIC) code OSIRIS [28, 43, 29], which can be applied to any massively parallel PIC code. On top of the parallel computing challenges presented by evolving hardware architectures, parallelizing a PIC code while maintaining load balance is inherently challenging on an algorithmic level. A PIC code contains two main data structures that comprise the computational load in a simulation: particles that can occupy any position in the simulation domain and field quantities that are discretized on a mesh grid. Parallelization is done by distributing particles and grid points among PEs, with load balance being achieved when the computational load, here defined as the calculation time, associated with these structures is distributed uniformly. OSIRIS [28, 43, 29] and most cutting-edge PIC codes use a spatial grid-based domain decomposition [40, 106, 107, 108], where each PE is responsible for a subset of the global spatial grid as well as any particles located there. While this domain decomposition algorithm is widely used, it suffers from the possibility that a large number of particles may move into a single PE (as is often the case for simulations of plasma-based acceleration, laser–solid interactions, or magnetic reconnection with pair production). Thus maintaining acceptable load balance can be challenging, given that the computational load will generally scale with the number of particles.

Various strategies have been implemented in an effort to dynamically load balance PIC codes using grid-based decomposition. Perhaps the most straightforward load balancing scheme is to enlarge distributed-memory spatial domains by using a large number of shared-memory cores on each PE and distribute the computational load evenly across these cores. This allows for localized load-imbalance situations, such as large density spikes, to be smoothed out, generally leading to good improvements in performance and scalability [43]. This offers limited relief, however, due to hardware limitations on the available number of cores. Another solution based on a static equidistant domain decomposition is known as the taskfarm alternative [109]. Here, each regular and equally sized distributed-memory sub-domain is further subdivided uniformly into tasks, with particles sorted accordingly. To process particles, each PE works serially through its own set of tasks before accessing, completing and returning

the tasks of other PEs with higher load. If tasks are small enough, load is balanced since no PE is ever idly waiting for remaining tasks to be completed. However, this approach can lead to complicated and costly communication patterns, especially if only a few PEs have a very large number of particles. A different approach is taken by Liewer and Decyk who pioneered the idea of shifting PE boundaries in their 1988 algorithm GCPIC [110, 40]. Computational load is projected onto one axis so that the problem of load balancing becomes one dimensional. Partitions are found along this axis, and the resulting domains are partitioned in the second and third dimensions in the same manner, resulting in rectangular (or cuboid, although we will use rectangular to describe both two- and three-dimensional domain shapes) partitions. In this approach, the corners of each rectangle may not aligned. Similar approaches based on rectilinear partitioning have been utilized both with aligned [43, 111] and unaligned [112] corners. Aligning corners simplifies communication patterns, but precludes obtaining a high degree of load balance for some problems. Dynamic load balancing can also be achieved by decomposing the simulation into many small units called tiles (or patches). Each PE handles one or more of these units, with the algorithm dynamically assigning them between PEs to maintain an even load [42, 106]. Similar strategies involving tiles have also been employed for effective parallelization on GPU architectures [41].

In this section, we describe an extension of the previous dynamic load balancing algorithm of Osiris [43]. In this new scheme the global simulation space is divided into many small, self-contained "tiles," which contain all particle and grid quantities for a particular region of space [113]. The parallel domain decomposition is determined by assigning one or more tiles to each PE such that computational load is balanced as much as possible; in addition, one or more threads are assigned to each tile within a given PE based on computational load. Thus the domain for each PE is constructed by using an irregular number of rectangular-shaped tiles whose corners are aligned. The ability to assign multiple threads to each tile—following the shared-memory parallelization algorithm already present in Osiris—allows us to significantly improve performance for simulations with small regions of high particle density by enabling

the parallel use of a multi-core PE on a single tile. This provides significant improvement over previous purely tile-based dynamic load balancing implementations [42, 106] that allow for only one thread per tile on CPUs. We find that this feature also allows for the use of larger tiles—reducing the overhead of passing particles between tiles—while still maintaining load balance. Our implementation also scales well with thread number (both when keeping the number of PEs times thread number fixed and when keeping the number of PEs fixed) and gives particularly large speedups for very imbalanced simulations, being well suited for efficient use in today's evolving HPC climate as available on-chip thread count continues to climb.

This section is organized as follows: In Sec. 2.1.2, we discuss the implementation of the tile structure into OSIRIS, including both distributed-memory and shared-memory parallelization schemes. We discuss the overhead and performance of the tiling scheme in Sec. 2.1.3 by analyzing simulations both with and without load imbalance, including a 3-D simulation of particle-driven plasma wakefield acceleration (PWFA). Compared to previous OSIRIS algorithms, our tile-based implementation of dynamic load balancing gives an order-of-magnitude increase in scalability with thread number and more than a factor of 2 overall speedup for two different physics simulations using 2048 cores. We expect even larger speedups are possible when using a larger number of computing cores.

2.1.2 Methodology

In the last two decades, the continued growth of top-ranked HPC systems has come almost entirely from increased parallelism, with present systems comprising up to $\sim 10^6$ cores. At the highest level of parallelism, these massive computer systems are viewed as a network of distributed-memory PEs amongst which the simulation can be partitioned. (Note that one can allocate multiple PEs within a single shared memory node; however, in this case memory is treated as distributed from a computing perspective.) Given that these PEs do not share memory, using a spatial domain decomposition requires that particle and field data

be exchanged between neighboring PEs at each time step. Domains for each PE should be structured such that the computation time on a given region is much larger than the time spent communicating boundary information, i.e., maximizing the computational volume to boundary surface area ratio to minimize parallelization overhead.

Previously, the domain decomposition in OSIRIS was structured such that each PE had only one neighboring PE in each direction, i.e., the corners of rectangular domains always matched up. The domain sizes could be statically assigned at the beginning of the simulation or changed dynamically throughout to maintain load balance [43]. Aligning domain corners simplified the communication pattern for the sharing of boundary information, but limited the achievable load balance. In an effort to improve dynamic load balancing and to enhance shared-memory parallelism in OSIRIS, we decompose the global simulation space into many small, static, regularly spaced rectangular regions called "tiles" which have aligned corners, following a strategy similar to Ref. [42, 106]. The domain decomposition is then determined by assigning a collection of one or more tiles to each PE such that computational load is balanced. These tiles can be exchanged dynamically between all PEs at chosen intervals throughout the simulation to maintain load balance.

2.1.2.1 Distributed-memory parallelism

To determine the most balanced parallel partition, i.e., assignment of tiles to PEs, we must first devise a way to quantify and organize the computational work required to process each tile. To this end, we create a load array with the same dimensions as the simulation and containing one entry per tile. Assuming that the computational load scales linearly with both the number of particles and the number of grid points,[1] we compute the array entry for the i^{th} tile, $\mathcal{L}_i$, as $\mathcal{L}_i = \mathcal{N}_{i,\text{part}} + \mathcal{C}N_{i,\text{cells}}$ as in Ref. [42], where $\mathcal{N}_{i,\text{part}}$ and $\mathcal{N}_{i,\text{cells}}$ are the number of particles and cells in the i^{th} tile, respectively, and the cell weight C represents the

[1] This is accurate for finite-difference-based field solvers. For other types of field solvers a different load formula can be straightforwardly derived.

computational load of a single cell compared to one particle. This depends on simulation parameters such as interpolation level or field solver type, and is left for the user to define. We found that for typical 2- and 3-D simulations $C \approx$ 0.5–2 is effective. The load array can then be partitioned into $\mathcal{N}_{\rm PE}$ regions of contiguous tiles such that the total load in each partition is as close as possible to the ideal average load, $\sum_i \mathcal{L}_i/\mathcal{N}_{\rm PE}$, where $\mathcal{N}_{\rm PE}$ is the total number of PEs. Partitioning a 1-D load array in this way is trivial since boundaries must be determined only along one dimension. However, partitioning a 2- or 3-D load array is more complex, and multiple solutions may be considered.

One of the approaches proposed by Saule et al. [112] (see Ref. [40] for an example of a PIC code using this implementation) is to consider 2-D partitions where each PE is left with a rectangular region of space, with each boundary connecting to one or more neighbors. Partitions where each boundary connects to a single neighbor (matching corners)—referred to as a rectilinear partition in [112]—was previously implemented in OSIRIS [43]. A more complex scheme, referred to as $P \times Q$ jagged in [112], allows for multiple neighbours along boundaries in the x dimension, and only one neighbor in y. This allows for decomposing the multi-dimensional load balance problem into separate uni-dimensional scenarios: we first define P partitions along x so that the load is evenly divided among them; we then proceed by dividing each of these P partitions into Q sub-partitions such that the computational load for all Q sub-partitions in a P partition is equal. This process can be straightforwardly extended to three dimensions. The $P \times Q$ jagged partition is intuitive and has rather simple boundaries between PEs, but for some computational load distributions it may not yield a perfectly balanced configuration, given that the boundary positions are limited to grid cell boundaries.

Another method to load balance in multiple dimensions is to place tiles sequentially along a space-filling curve [42]. It is important to recall that that physical domain in this method is broken up into tiles, where each tile is a rectangle with aligned corners. The challenge then is to assign a group of tiles to each PE. After estimating the optimal load per PE, we start with

the first PE and, following the curve over tiles, assign tiles to it until the load is close to the optimal value. We then proceed to the next PE/tile until all tiles have been assigned. This way the load imbalance can be adjusted on the granularity of a single tile. In reality, the load imbalance is more complicated because of differences in communication costs. The advantage of using a space-filling curve for ordering the tiles as opposed to, say, ordering tiles based on their Cartesian grid coordinates (as is done for MPI rank ordering), is that we will end up with simulation domains that are contiguous in space and have relatively simple boundaries with a small number of neighbors [114]. We implement the space filling curve methods by creating a 1-D load array, where tiles are ordered in the array using their position along the space-filling curve, then dividing this array into $\mathcal{N}_{\mathrm{PE}}$ segments with roughly equal load. Tiles falling on each segment will be assigned to the corresponding PE.

We implemented both the $P \times Q$ jagged partition and the space-filling curve methods in OSIRIS, with two choices for the space-filling curve. The first choice (Snake) simply passes through each tile by snaking back and forth in simulation space. This curve has no restrictions on domain size and yields the simplest boundaries, but does not maximize the ratio between computational volume and boundary surface area. The second choice (Hilbert) traces through the tiles using a Hilbert curve [115], which will maximize the ratio between computational volume and boundary surface area. However, this requires that the tile number be a power of 2 in the smallest dimension (and integer multiples of that number in other dimensions) and leads to more complex boundary shapes, potentially with more neighbors. Figure 2.1 shows schematics of various 2-D domain decompositions using 4 PEs for a simulation featuring a high-density diagonal stripe of particles (blue) surrounded by vacuum. We show a uniform partition without using tiles, as well as partitions using the three load balancing schemes implemented here, where the number of tiles is four times larger than the number of PEs. Note that with uniform partitioning, two PEs contain very few particles. For the density profile and tile size shown here, all three tile schemes achieve roughly the same degree of load balance across PEs. However, for more complicated density profiles the $P \times Q$ jagged

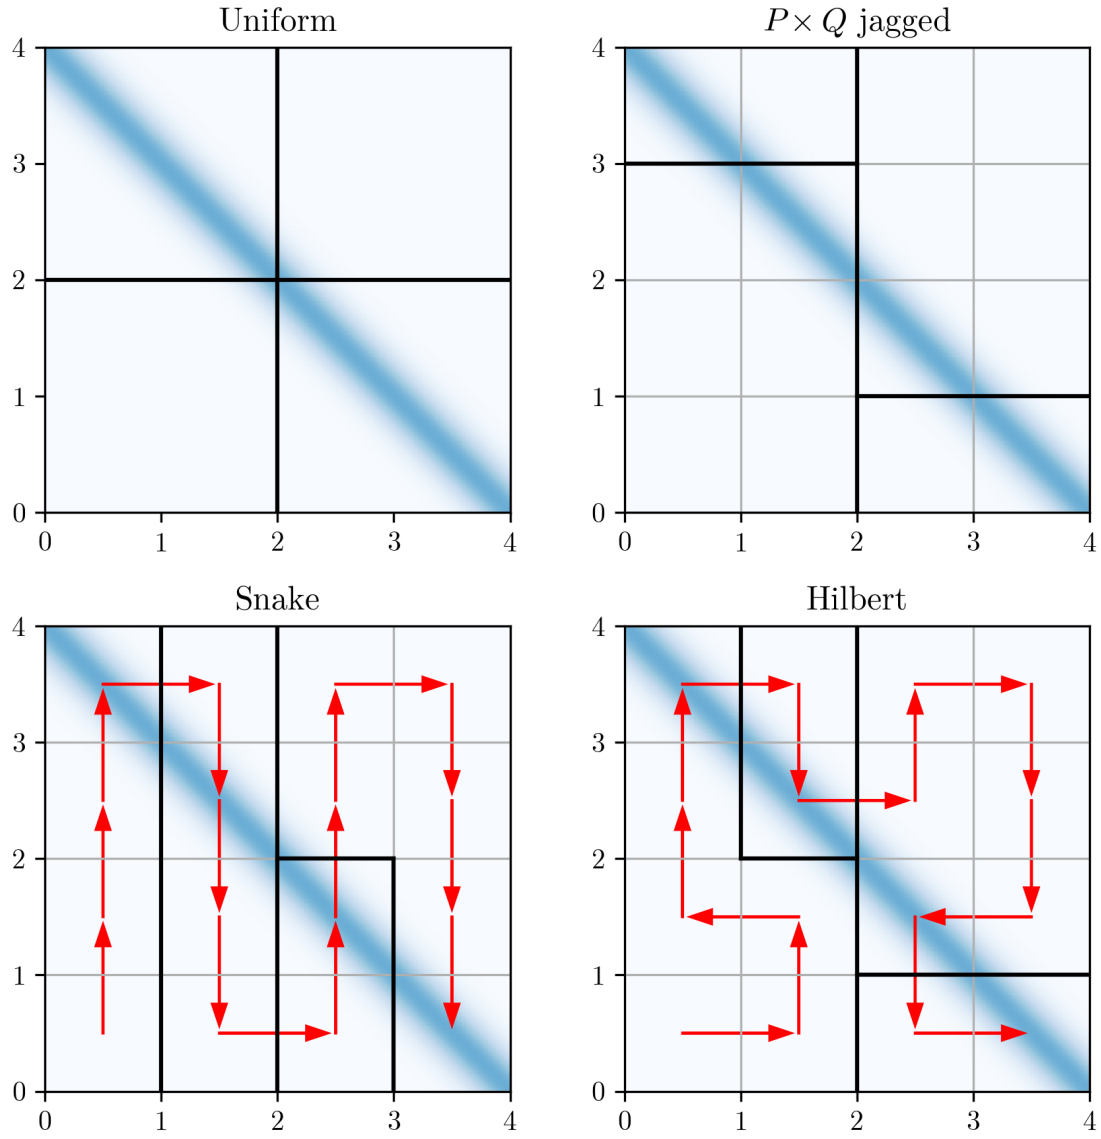

Figure 2.1: Domain decomposition using 4 PEs for a high-density diagonal stripe of particles (blue). We show a uniform partition without tiles, as well as partitions using the three load balancing schemes implemented here. Gray lines indicate tile boundaries, black lines indicate PE (i.e., MPI) boundaries, and red arrows trace out tile ordering along the space-filling curve. Note that with uniform partitioning, two PEs contain very few particles. For the density profile and tile size shown here, all three tiles schemes achieve roughly the same degree of load balance.

scheme will usually not load balance as well as the Snake/Hilbert schemes. To minimize communication overhead, we group messages from multiple tiles with the same destination PE into a single message, which greatly improves MPI performance.

To summarize, load balancing is performed by calculating the load array, determining the load-balanced domain decomposition, and then distributing the tiles among PEs accordingly. This can be done solely at the beginning of the simulation or dynamically throughout the simulation at chosen intervals using current simulation information. The domain decomposition defines a mapping between tiles and PEs, which is used to determine which tiles are to be sent and received, where tiles with the same destination PE are grouped into a single MPI buffer to minimize the effect of latency. Note that each successive domain decomposition is independent of the previous decomposition, i.e., a single PE may send all of its tiles to various other PEs and receive entirely new ones.

2.1.2.2 Shared-memory parallelism

The parallelization and load balancing algorithm described in the previous section works well for several problems, but does not guarantee ideal load balance for challenging scenarios where a large number of particles accumulate in a small number of tiles. Consider such a situation, where the load on a particular tile corresponds to α times the optimal load per PE, where $\alpha > 1$. In this case, the best the load balancing algorithm can do is to assign that tile to a single PE. However, the load for that PE will be α times the optimal value, and we will experience a slowdown by a factor of α. To address these situations and further improve parallel performance, we exploit the possibility for shared-memory parallelism that is available on most of today's computer systems, where a single PE may have many cores or threads that share memory with one another. While load balancing across distributed memory must be done with the coarse resolution of one tile, load balancing within a single PE with many threads can be done on a finer scale, for example, by dividing particles evenly among shared-memory threads.

For the case where multiple tiles are assigned to a single PE, there are two common approaches for processing the tiles using shared-memory parallelism. If there are many more tiles than available threads on the PE, load balance can be achieved by looping over tiles using a dynamic scheduler ("first come, first served"), with one thread per tile. This has been previously demonstrated by Ref. [106]. However, if one single tile contains a very large number of particles, the associated thread will take much longer than the others and load balance will fail.

Alternatively, tiles can be processed in serial with all threads working on a single tile at any given time. This way a single thread is never stuck on a tile with many particles. Dividing up particles in a single region of space amongst shared-memory resources is a strategy commonly implemented in GPU codes [41, 116]. However, given that all threads are assigned to the same region of space (tile), we must avoid memory collisions when doing the current deposition. This can be achieved either by using atomic operations, or by creating a separate electric current grid for each thread, then summing these arrays together after the current deposition is complete. Both options have specific drawbacks that impede parallel scalability for large numbers of threads, but the latter option generally gives better performance on CPUs and is used in OSIRIS [43].

We take a novel approach and include both types of shared-memory parallelism in a single framework. For a given PE with $\mathcal{T}$ threads and total load $\mathcal{L}$, let $\mathcal{L}_j$ be the load found on its j^{th} tile. Each time step we calculate which tiles should be processed in parallel (one thread per tile) and which tiles should be processed in serial ($\mathcal{T}$ threads per tile). A tile is deemed "heavy" if its load is greater than the average load per thread, i.e., if $\mathcal{L}_j \geq \mathcal{L}/\mathcal{T}$, and "light" otherwise. However, if a PE has fewer tiles than threads, all tiles are classified as "heavy" to avoid idle threads. For routines like those in particle processing, light tiles are processed first by assigning one thread per tile using a dynamic scheduler (we use OpenMP). Once all light tiles are completed, each heavy tile is executed one at a time, with all threads processing particles on that tile in parallel.

A schematic of the heavy-light tile organization is shown in Fig. 2.2 for a PE with five light tiles, two heavy tiles, and four threads. Using light tiles avoids data dependency issues when there are many more tiles than threads. The inclusion of heavy tiles ensures that no threads are idle in cases with (1) fewer tiles than threads on a PE or (2) disproportionately large load on a single tile. This implementation is critical for simulations with very high particle densities, such as plasma wakefield or shock simulations, where the distributed-memory load balance may result in PEs having just one tile containing a majority of particles. Note that tile G is queued for execution by the first available thread. If all light tiles have similar load, this leftover tile could double the execution time of light tiles compared to having only four light tiles. The number of light tiles could be adjusted to be divisible by the number of threads, or light tiles could be sorted by decreasing load to mitigate this overhead. However, we do not address this particular issue here since the impact is minimal.

Alternatively, it is possible to use a task-based approach to thread over tiles. The algorithm begins by creating one task per tile, for all tiles. Light tiles will then be processed without further parallelism, while heavy tiles will be processed using multiple threads per tile, with the number of threads being proportional to the number of particles in the tile. The number of threads for each heavy tile can be calculated as $\mathcal{T}_i = \lfloor \mathcal{T} \times (\mathcal{N}_i/\mathcal{N}) \rceil$ (round to nearest integer), where $\mathcal{T}$ is the total number of available threads in the PE, $\mathcal{N}$ is the total number of particles in the PE, and $\mathcal{N}_i$ is the number of particles in the tile. These tasks will be scheduled dynamically, assigning, if possible, a higher priority to light tiles. This approach should, in most scenarios, reduce the number of threads assigned to individual heavy tiles and ensure better load distribution across cores inside a PE, and will be further explored in a future publication.

2.1.2.3 Tile boundary management

When using a parallelization scheme based on spatial domain decomposition (such as the tiling scheme described here), each tile is required to exchange information with the neighboring

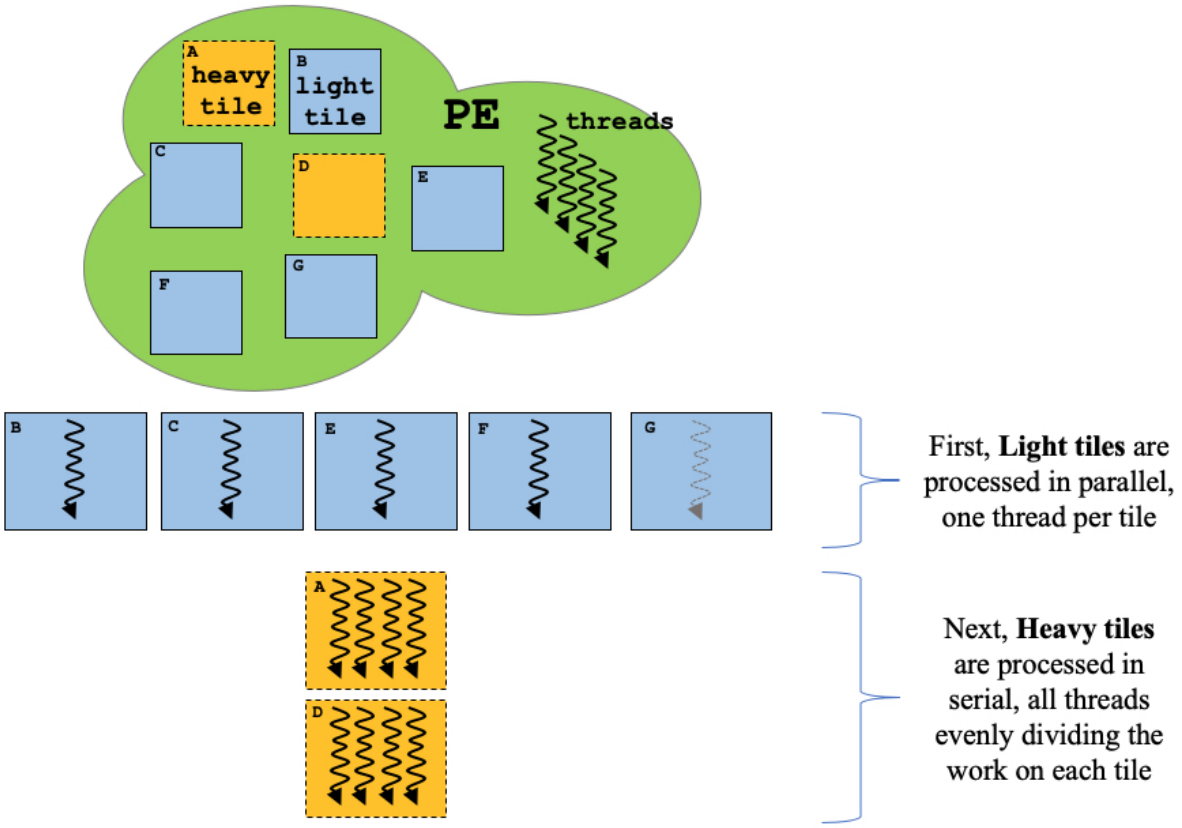

Figure 2.2: Work flow for shared-memory parallelization on a single MPI process with four threads. Each MPI process handles its tiles differently depending on their computational weight—a feature unique to our implementation. Using shared-memory threads (OpenMP in our case), tiles with load less than the average load per thread on that MPI process (light tiles) are processed in parallel with one thread per tile. The gray dashed thread in tile G represents a light tile queued for execution by the first available thread. Lastly, tiles with above-average load (heavy tiles) are processed in serial with all threads dividing work on that tile evenly.

tiles at each time step. This includes both grid information (edge values of fields and densities) and particle information (particles moving to/coming from another tile). Since each tile is self-contained, boundary information must be exchanged between all neighboring tiles—both tiles found within a single PE and those located across a PE–PE boundary. Note that by construction, tile corners are aligned. The exchange of boundary information introduces overhead not only from the MPI communications, but also from reassigning particles to different tiles within the same PE, with the latter overhead increasing as tile size decreases. If not optimized, these overheads can outweigh the benefit of the load balance algorithm, severely limiting its applicability.

In our implementation, boundary information between different PEs is packed into buffers and shared via MPI; tile boundary information within a single PE is instead referenced directly in shared memory. Whenever tiles are (re)assigned to PEs, each PE will gather and store information regarding which of its tile boundaries are internal and which are external. Each time step when boundary information is exchanged, boundaries are processed sequentially over the number of dimensions to avoid corner (diagonal) communications. When exchanging grid information, we loop over all tiles using parallel threads; local tile boundary values are copied directly in shared memory (no buffer required), and boundary values to be sent to other PEs are packed into one buffer per target PE and sent via MPI. We then loop over all tiles with external boundaries (again using parallel threads) to copy in boundary values from the received communication buffers. Particles moving between tiles are processed in a similar manner. Each tile maintains buffers to store exiting particles; buffered particles moving to a local tile are referenced and unpacked, while particles moving to a different PE are packed and sent over MPI. Grouping MPI messages going from a local PE to the same target PE greatly improves performance (as opposed to using one message per tile boundary, for example) by reducing the number of messages and limiting the impact of communication latency, since message sizes are larger.

Using the tiling scheme with either of the space-filling curves will lead to a small increase

in communication time when compared to a uniform partition. This is mainly due to an increased number of neighboring PEs (the corners of each PE are no longer aligned) that communicate smaller messages (i.e., communication is mostly latency-dominated). This increase in communication time can be offset by obviating the particle sorting step required in most PIC codes while still ensuring data locality and cache coherency (due to the small size of a tile). Furthermore, the benefit of proper parallel load balance can greatly outweigh any penalties incurred from communication overhead for otherwise unbalanced runs.

2.1.2.4 Best practices

In this section we briefly discuss parameters a user or developer should consider when working with a tile-based PIC code. First, the choice of tile size has a significant impact on the performance of the algorithm: smaller tile sizes improve data locality and load balance at the expense of a higher overhead from tile boundary management, as the number of particles (compared to the total particles in a tile) crossing boundaries will increase, as will the number of boundaries between PEs. On the other hand, larger tile sizes reduce the boundary management overhead but can interfere with data locality and hinder parallel load balancing by limiting the available resolution of the parallel partition. The optimal tile size will be problem-specific, as discussed in further detail in Sec. 2.1.3.1.

Second, the choice of how frequently to perform the dynamic load balance is problem-specific, but in our experience it should be done rather frequently, e.g., once every 20 or 40 time steps. For all simulations in Sec. 2.1.3, the amount of total simulation time spent performing the dynamic load balance was always below 4%, (usually between 1–2%). Good load balance was maintained for a very rapidly evolving plasma with little overhead by repartitioning every 20 time steps.

Finally, perhaps the biggest concern for developers and users of a tile-based PIC code is memory management. Specifically, extra buffers need to be maintained to store particles migrating between tiles and PEs. As mentioned briefly in Sec. 2.1.2.3, our OSIRIS implemen-

tation maintains two types of buffers at the tile level and one extra type of particle buffer at the PE level. Each tile must have (1) a buffer to hold all of its particles and (2) a buffer (or multiple) to store the particles that have left its domain and need to be moved elsewhere. In addition, each PE requires separate buffers to store particles to be communicated to various neighbor PEs. The sizes of buffers used to store migrating particles must be chosen carefully to maximize performance (e.g., limit reallocation and number of MPI messages) without running out of memory. When buffers are grown to include additional particles, increasing the buffer size by only a small amount may trigger memory reallocation every few time steps during a large influx of particles. However, if buffers are grown by a large amount, a small influx of particles across a large portion of the simulation space could consume a large amount of memory. In our case, we give the user the ability to define the initial size of the main particle buffers in each tile, along with the amount by which the buffer is grown when necessary. These parameters are problem-specific, and may require some experimentation to appropriately define. Internal memory management could also be used to write checkpoint data when the allocated memory approaches the hardware limit.

2.1.3 Results

To evaluate the performance of our dynamic load balance algorithm using tiles we will benchmark it against the baseline performance of OSIRIS using a static, regular spatial domain decomposition (referred to as "no tiles" or "OSIRIS" in all figures). As mentioned earlier, these results depend on the type of problem and level of imbalance, as well as user choices such as tile size and dynamic load balance frequency. We will analyze simulations of a uniform warm plasma, an ambipolar diffusion problem, and a plasma wakefield accelerator. All simulations were performed on Haswell compute nodes of the Cori system at NERSC, each with two sockets, and each socket containing a 2.3 GHz 16-core Haswell CPU (Intel Xeon Processor E5-2698 v3) supporting 2 hyper-threads, leading to a total of 64 hardware threads per node.

2.1.3.1 Uniform warm plasma

A uniform warm plasma is a perfectly balanced problem that can be simulated using standard static spatial domain decomposition techniques with excellent parallel efficiency, since the load for every PE is uniform throughout the simulation. Simulating a warm plasma is thus ideal for benchmarking the overhead of the tiling algorithm compared to the default parallelization strategies—both pure MPI and MPI/OpenMP hybrid parallelization options are available in OSIRIS (see Ref. [43] for details)—and optimizing tile size for best performance. The default MPI/OpenMP hybrid parallelization scheme divides the simulation space evenly amongst PEs, then processes particles within an entire PE domain in parallel amongst threads, with each thread using a separate electric current grid that must be summed at each time step (see Sec. 2.1.2.2). We simulate a uniform warm plasma in 2D and 3D using four compute nodes, for a total of 128 cores (256 threads). The product of the number of MPI processes (M) and the number of threads per process (N) is kept constant at $M \times N = 256$, for $N = 2$, 4, 8, 16 and 32. The 2-D runs use a 1024^2 cell grid with cell size $(0.0174\,\mathrm{c}/\omega_p)^2$ and a total of ~ 150 million particles, and the 3-D runs use a 128^3 cell grid with cell size $(0.0211\mathrm{c}/\omega_p)^3$ and a total of ~ 134 million particles; both runs use a time step of $0.012\,\omega_p^{-1}$. Particle velocities were initialized from a thermal distribution with a proper thermal velocity of $u_{\mathrm{th}} = 0.1\,\mathrm{c}$. All simulations were done using second-order (quadratic) interpolation for particles with periodic boundary conditions in all directions, and were run for 1000 time steps. Tile size was 16^2 cells in 2D and 8^3 cells in 3D. Figure 2.3 shows the performance of the code for the various cases.

The performance of the default hybrid parallelization in OSIRIS is shown to decrease with increasing thread number (decreasing PE number). This is to be expected as the shared-memory algorithm is required to do additional work (zeroing additional current arrays and reducing the results from all threads) with increasing thread count. Since all tiled simulations in this configuration have only "light" tiles (16 tiles for every thread), the new tiling algorithm does not suffer from this limitation and shows good scalability all the way up to 32 threads. The drop in performance is mostly due to the fact that not all routines

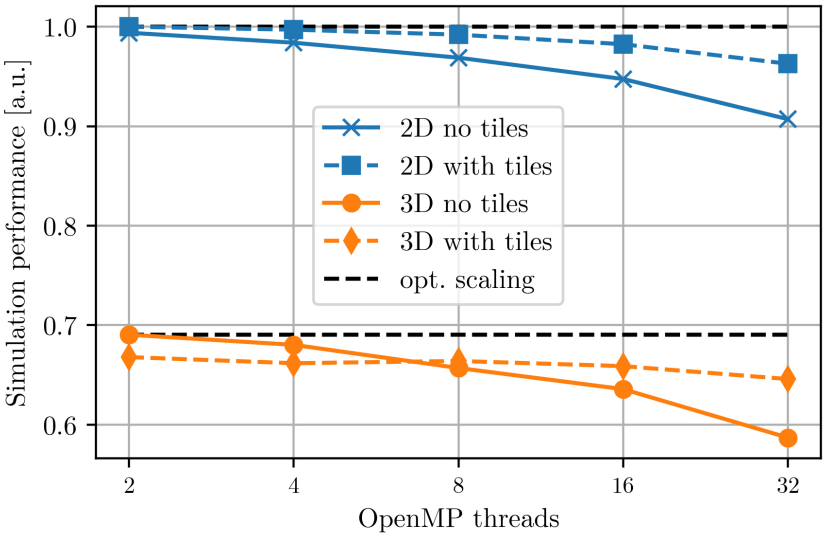

Figure 2.3: Thermal plasma simulation performance in 2D and 3D for varying number of threads, with (dashed line) and without (solid line) tiling algorithm. Tiles were 16 cells square in 2D and 8 cells cube in 3D. Number of MPI processes times number of threads per process was kept constant at 256. Performance is calculated as the number of particles pushed (including field solve and other elements) per second, normalized to the fastest run.

have been shared-memory parallelized, as they do not represent a significant overhead for small thread counts.

The new algorithm is only outperformed by default OSIRIS for 3D geometry with small (2,4) number of threads. This is related to the overhead of moving particles between domains; default OSIRIS without tiles uses larger domains, so there will be fewer particles (compared to the particles in the domain) crossing boundaries. However, for larger thread counts the limitations of default OSIRIS outweigh this overhead, and the new algorithm performs much better.

To determine the optimal tile size, we repeated the above simulations for varied tile sizes and number of particles per cell. The latter parameter is important because it directly impacts the computational load of each tile. The number of particles per cell were varied between 1, 16, 64 and 144, and the tile sizes were varied by powers of 2, with values ranging from 8–64 cells on each side. Smaller values are not possible using second-order (quadratic) interpolation, as a minimum of 5 cells is required. Table 2.1 summarizes our results. We found that the optimal tile side length varied between 16–64 cells in 2D and 8–16 cells in 3D. These results are consistent with our expectations: when the amount of computation per tile is larger (as is the case with higher numbers of particles per cell), the overhead of moving particles between tiles has a smaller impact, and the benefits of smaller tiles in terms of data locality lead to higher performance. For smaller numbers of particles per cell this is no longer the case, and larger tiles perform better.

2.1.3.2 Ambipolar diffusion

Ambipolar diffusion is of particular importance to a large range of physics scenarios, such as laser–solid interactions and inertial confinement fusion [117, 5, 118]. This is a particularly difficult problem to simulate using spatial domain decomposition as the plasma, which is initially confined to a small region of space, expands into vacuum (or near vacuum). We will use this problem to test the effectiveness of our algorithm under extreme load imbalance.

	2D			3D	
Particles per cell	1	16	64+	1	8+
Optimal tile size	32–64	32	16	16	8

Table 2.1: Optimal tile sizes (cells on each side) in 2D and 3D for a thermal plasma with varied particles per cell. Decreasing the tile size introduces more overhead but improves cache locality, which is important for simulations with many particles per cell.

We perform a simulation of the expansion of an electron-ion plasma undergoing ambipolar diffusion in 3D, using the various parallelization strategies discussed in Sec. 2.1.2. We start with a constant-density sphere of electrons and ions, where the electrons have a temperature of 130 keV, the ions are cold, the ion-to-electron mass ratio is 1836, and each species has 1000 particles per cell. The sphere of particles has radius 1.4 c/ω_p and is stationed in the center of a cubic periodic box of side length 16 c/ω_p with 160^3 cells. The electrons quickly diffuse outward, but are soon slowed by the ambipolar space charge fields of the ions. The ions are in turn pulled out and start to expand, which slows down the expansion of the electrons. Afterwards the electrons will again start to diffuse, leading to an oscillatory expansion of both species. We ran each simulation a total time of 300 ω_p^{-1} with a time step of 0.0577 ω_p^{-1}, and we used first-order (linear) interpolation for particles.

Figure 2.4 shows the combined electron-ion particle density for an analogous 2-D simulation at times 0, 153, and 300 ω_p^{-1} after the beginning of the expansion, where ω_p is the electron plasma frequency before expansion. The domains of the 16 PEs overlay the density; density slices along any dimension in 3D show similar behavior, but the parallel domain decomposition can be better visualized and understood in 2D. For the first half of the 3-D simulation, about 70% of the particles are contained within a sphere of radius 2.5 c/ω_p, or just 1.6% of the entire simulation volume. By the end of the simulation, 70% of the particles are contained

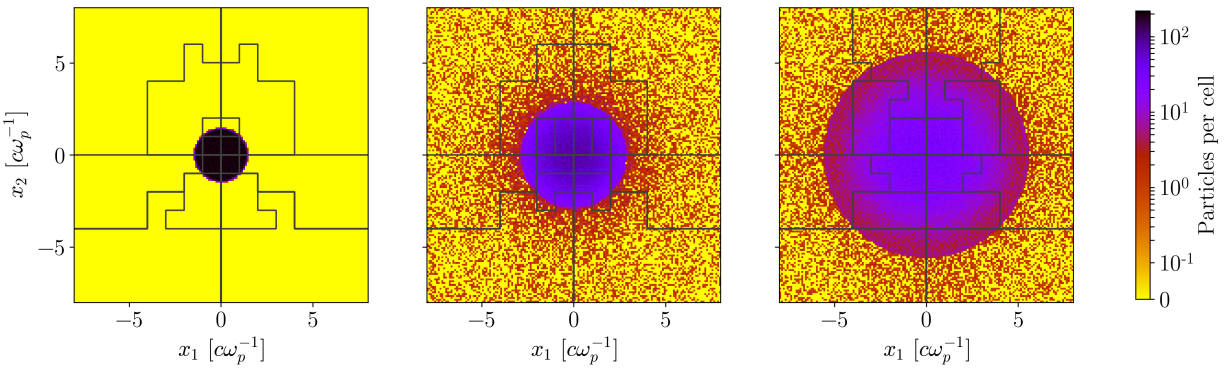

Figure 2.4: The 16 PE subdivisions overlaying particle density of the ambipolar diffusion problem in 2D at various times. Cell weight is 2.0 for this case.

within a sphere of radius 5 c/ω_p (13% of the total volume). Since the particles largely reside in the center of the simulation box, a traditional spatial domain decomposition with uniform partition sizes will only allow good parallel load balance for up to 2 PEs per dimension (8 PEs total), and will show severe imbalance if the number of PEs is increased to 4 or more in any dimension. This simulation is also challenging for our tile-based approach: using a total of 4096 cube-shaped tiles with 10 cells to a side, 70% of the particles are contained within just 56 tiles for the first third of the simulation, steadily increasing to 432 tiles by the end of the simulation.

We test the strong scaling parallel speedup of this simulation, keeping the problem size constant and running on 128 to 2048 cores for various cases: the default hybrid MPI/OpenMP algorithm without ("Osiris, no dlb") and with ("Osiris dlb") the previously implemented dynamic load balance [43], light tiles only and heavy tiles only, each with dynamic load balance using the Hilbert space-filling curve ("light tiles" and "heavy tiles", respectively), and a combination of light/heavy tiles with dynamic load balancing using all three schemes ("Hilbert," "Snake," and "$P \times Q$ jagged"). Recall that the previously implemented dynamic load balance uses the rectilinear partition from [112], similar to $P \times Q$ jagged but with only one neighboring PE on each domain side. As mentioned in Sec. 2.1.2.2, the shared-memory parallelization of light and heavy tiles is as follows: light tiles are assigned only one thread

each, with tiles processed in parallel on a first-come-first-served basis; heavy tiles are processed in serial, with all threads working on a single tile at once. The two cases without tiles use the maximum number of threads per MPI process without hyperthreading (16), as this is most favorable for the default hybrid algorithm. The cases with tiles use 2, 4, 8, 16, and 16 threads per MPI process for the five runs with increasing core counts, respectively. Dynamic load balancing was performed at every 20 time steps and occupied just a few percent of the total simulation time. The chosen cell weight parameter was $C = 2$. For the "Hilbert" run using 2048 cores, dynamically load balancing every 10 and 40 time steps caused the total run time to increase by 10% and 11%, respectively. Figure 2.5 shows the results in terms of (a) overall simulation performance and (b) load per core, normalized to the fastest simulation with 128 cores.

The default hybrid algorithm without dynamic load balancing is unable to scale due to load imbalance issues. Figure 2.5 also shows that the scalability of the light tiles algorithm is severely limited as a result of its inability to parallelize within each tile, and in fact the light tiles algorithm is the slowest algorithm at all core counts. To explain the poor performance when using only light tiles, consider the simulations with 1024 or 2048 total cores, where each PE has 16 threads. During the first third of the simulation, the 56 computationally expensive tiles could be spread across just 4 PEs (64 cores) since the algorithm groups together tiles close in space and requires that each PE has at least as many tiles as threads. Each light tile can only be processed by a single core, leading to significant performance degradation that worsens with increasing core counts. Using only heavy tiles significantly boosts performance since a PE may be assigned just a single tile, but note the drastic drop in performance for the first run using 16 threads (1024 total cores). This overhead is a result of summing the 16 separate electric current grids for each tile, which is especially detrimental for PEs with many tiles containing few particles.

The limitation in parallel scalability when using only light or only heavy tiles is overcome by our new algorithm that uses a combination of light and heavy tiles. We implement an

additional level of parallelization within each tile, which allows a PE to efficiently devote all of its cores to process a single computationally expensive tile. As seen in Fig. 2.5, our algorithm (a) maintains much better overall parallel scalability compared to any other method and (b) achieves near ideal load balance for up 1024 cores. At all core counts, the combination of heavy/light tiles always outperforms the use of light tiles only, being ~ 2 times faster for 128 cores and ~ 6 times faster for the largest core count. The average number of heavy tiles over the course of the simulation steadily increases with thread count, ranging between 0.5% and 8.2% of all tiles from 128 to 2048 cores, respectively. The small drop in (b) from ideal scalability at 2048 cores has to do with the problem/tile size: at this core count we will have on average only 2 tiles per PE, which does not allow for effective load balancing. Reducing the tile size may improve performance for this case.

When comparing the various load-balancing schemes, using the Hilbert space-filling curve was consistently the fastest, and with 2048 cores it was 1.2, 1.5, and 1.4 times faster than the $P \times Q$ jagged scheme, snake space-filling curve scheme, and previous dynamic load balance, respectively. An overall speedup of a factor of 2.2 was gained compared to original OSIRIS without dynamic load balance. Although the snake space-filling curve consistently gave excellent load balance, extra overhead from MPI communications due to the shapes of PE boundaries ultimately caused those runs to be slower than the other topologies. For these problems there were only an average of $\sim 11,000$ particles per core for the case with 2048 cores, and many of the tiles have very few particles.

2.1.3.3 Plasma wakefield acceleration

Since its inception in the seminal paper by Tajima and Dawson [52], the field of plasma wakefield acceleration has been an active area of intense research. Plasma based acceleration (PBA) [119] is an accelerator scheme which uses an intense electron bunch—plasma wakefield acceleration (PWFA) [60]—or a laser—laser wakefield acceleration (LWFA)—to accelerate particles. As the beam propagates through the plasma, the electrons are pushed away from

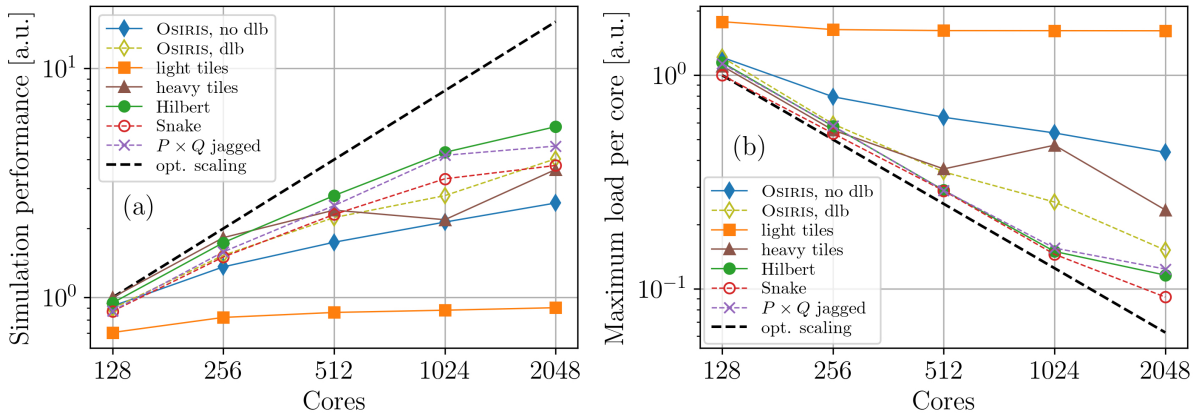

Figure 2.5: Strong scaling test of a 3-D ambipolar diffusion simulation. Performance in (a) is calculated as the number of particles pushed (including field solve and other elements) per second, and load in (b) is calculated as the maximum time any core spent in the advance-deposit routine, both normalized to the fastest run with 128 total cores. The "OSIRIS, no dlb" and "OSIRIS, dlb" runs are both done without tiles, one without load balancing and one with the previously implemented dynamic load balance. The "light tiles" and "heavy tiles" runs use the new dynamic load balance via the Hilbert space-filling curve; a light-tile run processes one or more tiles per thread for all tiles, while a heavy-tile run uses all available threads to process each tile sequentially. The "Hilbert," "Snake," and "$P \times Q$ jagged" runs use the new dynamic load balance along with a combination of light and heavy tiles, where the names refer to the load-balancing scheme used. The two cases without tiles are performed with 16 threads per MPI process, and tiled runs use an optimal configuration for each run (2, 4, 8, 16, and 16 threads per MPI process for the five runs with increasing core counts, respectively).

the beam to leave an exposed ion column, resulting in a region of space supporting large electric fields that travels at nearly the speed of light. Particles can be externally injected or self-injected (trapped) from the background plasma through various mechanisms [67, 120, 121, 122, 123, 9], then accelerated to high energies. The excitation of the wake by the drive beam is strongly nonlinear, as can be the evolution of the driver and the injected and accelerated beams, making PIC simulations the tool of choice for modeling these scenarios.

Simulations of PWFA naturally contain regions of very high particle density—the injected particle bunch and beam driver—that are dynamic in nature, surrounded by regions of relatively low-density background plasma. These density distributions can result in severe load imbalance, for which finding a single PE decomposition that balances load for the entire simulation is nearly impossible. Effective dynamic load balancing of the simulation is therefore crucial for efficient numerical modeling of PWFA, particularly when using large core counts. We test the performance of our dynamic load balancing algorithm with a PWFA scenario similar to that studied by Dalichaouch et al. [124]. The driving beam is initialized with a radius of $\sigma_r = 2.1\,\mathrm{c}/\omega_p$ and then evolves self-consistently, focusing down to a radius of $\sigma_r = 0.2\,\mathrm{c}/\omega_p$ in the moving simulation window. When initializing a beam, to provide good statistics we use a fixed number of particles per cell and variable weights on the particles to vary the density. The beam evolution leads to a large load imbalance due to the large number of simulation particles concentrated in a small cell volume. Additionally, the accelerating structure formed with particles from the background plasma will show a density spike at the back of this structure that can be several orders of magnitude larger than the background density, creating a second load imbalance region.

To illustrate the complexity of finding an optimal domain decomposition, we show in Fig. 2.6 a snapshot from an analogous simulation using the "quasi-3D" version of OSIRIS that uses a cylindrical geometry (two-dimensional r-z in space, azimuthal expansion in θ) [125, 30]. Load balance was performed using the Hilbert space-filling curve and a cell weight $C = 1$ for 2048 PEs; the simulation size was 2048×1536 cells, with 256×128 tiles each of size

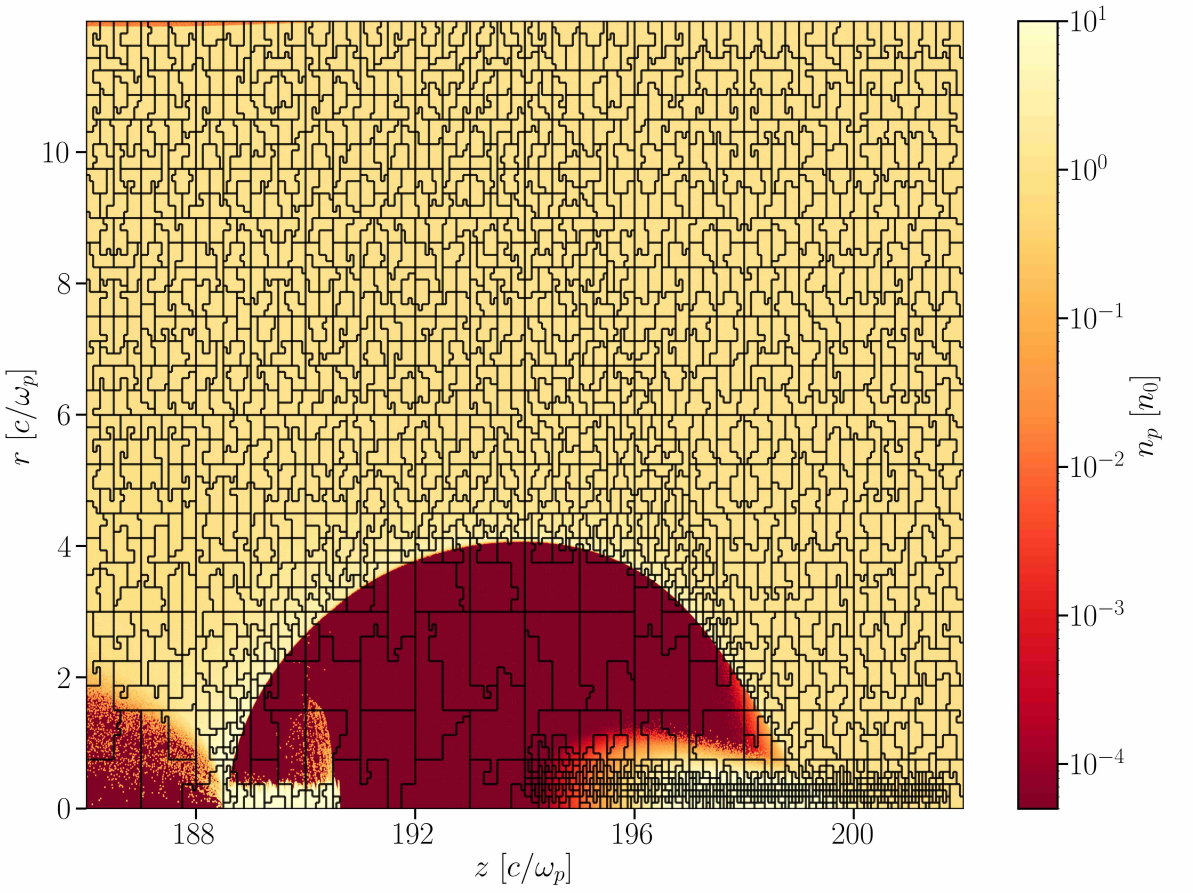

Figure 2.6: MPI subdivisions (black lines) overlaying electron density for a quasi-3D PWFA simulation with 2048 MPI processes.

8×12 cells. The total simulation space was of size $16 \times 12\,\mathrm{c}/\omega_p$, with 8 particles per cell for each of the driver and background electron species, and we used second-order (quadratic) interpolation for particles. This snapshot is about halfway through the simulation when the particle driver has focused tightly near to the axis, and the figure shows an overlay of PE boundaries with electron density. Note the large concentration of small domains near the beam driver, bubble sheath and injected particles, and the large domains in the near-vacuum region. Similar behavior occurs for 3D simulations, but the visualization of these structures is cluttered and brings little insight.

We tested the efficiency of our algorithm on 3D simulations of this PWFA scenario by performing a set of simulations with different parallelization options. All simulations were performed on 64 compute nodes with the architecture described above (2048 cores total) and varying numbers of MPI processes/threads, for the default hybrid parallelization without load balance and for the dynamic light/heavy tiles parallelization. The simulation size was $512 \times 768 \times 768$ cells, and the tile size was $16 \times 12 \times 12$ cells for the tile-based simulations. The total simulation space was of size $16 \times 24 \times 24\,\mathrm{c}/\omega_p$, with 1 particle per cell for each of the driver and background electron species and second-order interpolation for particles. The simulation was run for a total of 12800 time steps to a time of 200 ω_p^{-1}, with dynamic load balancing performed at every 40 time steps. For load calculations we chose a cell weight parameter of $C = 0.5$.

Figure 2.7 shows the evolution of the parallel particle load imbalance over the course of a simulation using 2048 total cores with two MPI/OpenMP configurations (see Table 2.2 for more detail) as the ratio between the maximum computational load on a single PE and the average computational load across all PEs. The load imbalance is shown for both the static uniform partition (no tiles) and the dynamic partition (tiles). As shown in the plot, the high-density regions near the propagation axis lead to a severe load imbalance about halfway through the simulation. For 2048 processes each with 1 thread, the imbalance reaches a peak (average) value of 37.1 (10.6) without dynamic load balancing. This result can be

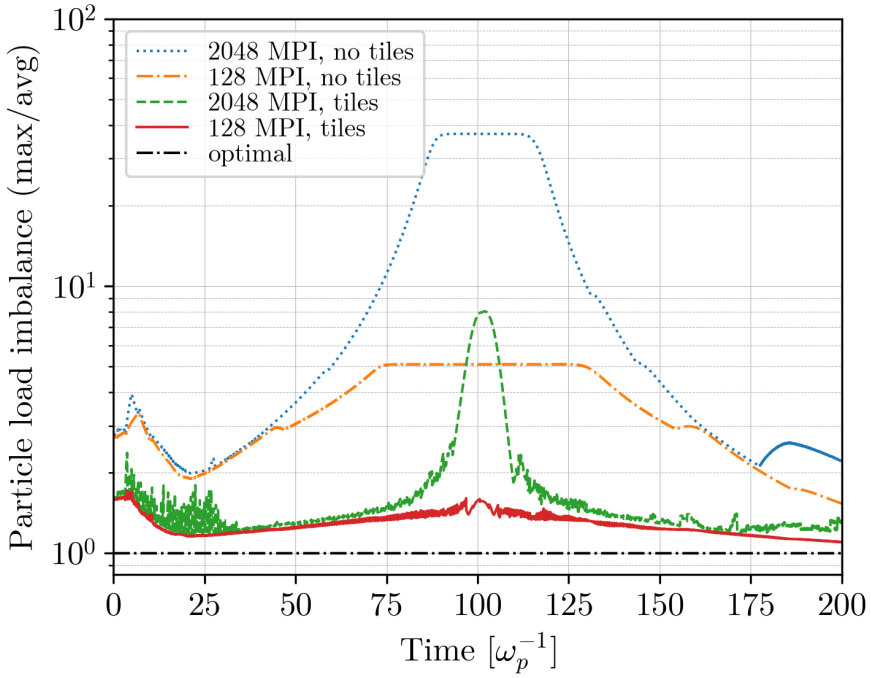

Figure 2.7: Particle load imbalance (maximum/average load) of a 3-D PWFA simulation using 2048 total cores with and without dynamic load balancing using tiles for two different MPI/OpenMP configurations. The dot-dashed line shows an ideal load imbalance of 1.0.

improved through the use of the hybrid MPI/OpenMP algorithm [43] that smears out the load imbalance by allowing for large domains to be assigned to each PE, leading to a peak (average) imbalance value of 5.1 (3.5) at 128 processes each with 16 threads. It should, however, be noted that this large number of threads leads to some performance degradation due to overhead, as described in Secs. 2.1.2.2 and 2.1.3.1. To further reduce the parallel load imbalance we must use the dynamic tile load balance. Using this algorithm we can lower the peak (average) value of the imbalance down to 8.0 (1.7) with 2048 processes, as shown in Fig. 2.7. Furthermore, using our heavy/light tile parallelization we can assign multiple threads to the heavy tiles, achieving a peak (average) load imbalance of just 1.7 (1.3) with 128 processes. Note that for this run, nearly all tiles were processed as light except for an average of 175 heavy tiles near the focus of the beam driver during the middle 10% of the simulation time.

Finally, we compare the total simulation times for the different configurations of MPI processes and OpenMP threads for the static uniform partition and the dynamic tiles partition. Table 2.2 shows the results, with time values normalized to the fastest simulation (512 MPI processes with 4 OpenMP threads and dynamic load balance). As shown in the table, performance of the static partition (no tiles) can be improved by over a factor of 2 simply by increasing the number of threads. However, the tile-based algorithm maintains good performance for all thread counts and always outperforms the static partition, being 2.18 times faster between the fastest runs. Cumulative load balance is roughly the same for tile-based runs with 4, 8, and 16 OpenMP threads, but the overhead discussed in Sec. 2.1.2.3 is slightly larger with 16 threads. The overhead associated with the dynamic load balance effort, however, was found to be only $\sim 1\%$ of the total simulation time in all cases. We also see that the timings for the tile-based algorithm vary less than 10% with varying number of threads, which is to be expected, as the computation should scale well with number of threads. This small variation highlights the versatility of the algorithm in terms of efficiency on various architectures supporting different numbers of threads. We note that for simulations with

MPI processes	OpenMP threads	Time [a.u.] no tiles	Time [a.u.] tiles
2048	1	4.63	1.19
1024	2	4.55	1.03
512	4	2.45	1.00
256	8	2.57	1.03
128	16	2.18	1.10

Table 2.2: Total simulation time with and without tiles (and dynamic load balancing) for the 3-D PWFA simulation. The total number of cores was kept constant (2048), only varying the number of OpenMP threads and MPI processes.

many more particles per cell, the speedup of the total simulation will more closely mirror the improvement in particle load balance. A more systematic study will be undertaken in the future.

2.1.4 Summary

Many of today's frontier problems in plasma physics necessitate fully kinetic simulations, which in many cases can only be achieved through particle-in-cell simulations (this is the case for the physics problems addressed in this dissertation). Although the PIC algorithm has been quite successful in addressing many kinetic plasma problems (and can be extended straightforwardly with additional physics models), it is numerically expensive, often requiring large-scale, massively-parallel computational resources. However, the traditional paralleliza-tion of the PIC algorithm is susceptible to parallel load imbalance. In problems such as plasma-based acceleration, laser–solid interactions, and magnetic reconnection with pair production, simulation particles may accumulate in small regions of space, leading to an uneven computational load. Dynamic load balancing across distributed-memory processing elements (PEs) and the efficient use of shared-memory cores are therefore essential to perform

large-scale simulations of these physics problems.

In Sec. 2.1 we have presented a novel hybrid parallelization strategy for the PIC algorithm that combines two different shared-memory parallelization algorithms, achieving excellent performance even for simulations with extreme imbalance. Our algorithm uses small, regularly spaced, self-contained regions of the simulation space that we refer to as tiles. These tiles contain all particle and grid quantities for a particular region of space and are dynamically traded between all PEs at chosen intervals to maintain computational load balance. Unique to our tile-based implementation of dynamic load balancing is the ability to assign either one or more than one thread to each tile depending on computational weight. This is especially useful for simulations with localized high-density regions for which the load-balanced domain decomposition may result in a PE having a single tile containing a majority of particles, or for simulations where a given PE has fewer tiles than threads. Furthermore, the ability to assign multiple threads to each tile allows us to use larger tiles, which can reduce the overhead of the tiling algorithm while still maintaining good load balance and high performance.

Our algorithm was shown to perform well for balanced simulations: it was on par with or faster than the traditional parallelization algorithm on the same hardware and scaled better with increasing thread count, with an overhead of less than $\sim$5% for the highest thread count. Our algorithm gave a speedup of more than a factor of 2 compared to simulations without dynamic load balance of an expanding plasma undergoing ambipolar diffusion, a critically difficult problem to parallelize with 70% of the particles contained within only 13% of the simulation volume. In particular, it was $\sim$ 6 times faster than running the simulation using only 1 thread per tile and attained near-ideal load balance for 8 times as many cores compared to running without tiles. We also analyzed the performance of our algorithm on plasma accelerator simulations and verified a speedup of over a factor of 2 when compared to performance without dynamic load balance, with the dynamic load balance itself only taking about 1% of the total simulation time. Even greater speedups are expected for larger simulations, where shared-memory thread number is small compared to

total computing resources. Our results also show that the performance of the algorithm does not vary significantly with different combinations of PE and thread numbers.

This algorithm was tested on a CPU architecture using MPI/OpenMP for distributed/shared memory parallelism, and can be straightforwardly extended to other architectures and programming models, such as the MPI-3 Shared Memory model, Coarray Fortran, extension to GPUs using, for example, CUDA, or extension to other architectures using Intel oneAPI. We are currently investigating the best way to implement this load-balancing algorithm on many-core accelerators, noting that similar algorithms have been shown to operate efficiently on such architectures [126, 42]. The software may need to be carefully structured to ensure that all threads are executing the same kernel while allowing for a combination of heavy/light tiles. The algorithm can also be combined with vectorized versions of the PIC algorithm, efficiently exploiting all parallelism levels available in present and near-future HPC systems and opening new avenues for the numerical simulation of kinetic plasmas.

2.2 Increased fidelity for the quasi-3D algorithm

For many problems of interest in plasma-based acceleration and laser–plasma interactions of short laser pulses, there is near-azimuthal symmetry. Pure r-z, two-dimensional algorithms cannot model a linearly (or circularly) polarized laser; they can, however, model radially or cylindrically polarized lasers. To circumvent this challenge, Lifschitz et al. [125] proposed an algorithm that decomposed (expanded) all fields (electromagnetic fields and current densities) into azimuthal harmonics. They noted that a linearly polarized laser with a symmetric spot size contains only the first-order azimuthal harmonic. Therefore, by keeping only the first azimuthal harmonic, laser–plasma interactions for lasers with symmetric spot sizes can effectively be studied in two dimensions while retaining the correct three-dimensional geometric effects. The particle locations in ϕ are updated such that the code is hybrid: it uses a PIC description in r-z and a gridless description in ϕ. The expansion can be truncated

at a small number of harmonics if near-azimuthal symmetry exists, thereby making the computational cost similar to a two-dimensional code. This algorithm was implemented into OSIRIS as described in Ref. [30], and the algorithm was coined as a quasi-3D description.

The OSIRIS implementation [30] differs from the original work of Lifschitz et al. [125] in two respects. The first is that it described a charge-conserving current method that is compatible with the quasi-3D description, and the second is that the field discretization was shifted by a half-cell to align with cylindrical routines previously in place in OSIRIS. Furthermore, as it was fully integrated into OSIRIS, it was able to take advantage of other optimization and physics features. Other PIC codes, such as SMILEI [106], that have implemented a finite-difference version of the quasi-3D algorithm use the field discretization proposed by Lifschitz. The discretization in OSIRIS appears to present some challenges, specifically that there can be a spurious growth in the E_r field right at the $r = 0$ axis. It is unclear whether similar field growth is observed in the other discretization, or if this problem is unique to the cell arrangement in OSIRIS. We note that a spectral-based quasi-3D algorithm was developed by Lehe et al. [127]; however, such an algorithm does not scale as well on massively parallel compute systems.

We have used quasi-3D simulations extensively for the work to be described in Chapter 3 for studying self-modulated laser wakefield acceleration (SM-LWFA) combined with directed laser acceleration (DLA). To enable these simulations, we needed to address some subtle numerical issues. In this section we describe this effort.

We first propose an improvement to the charge weighting of particles that originate near the axis that eliminates an on-axis density spike and corresponding field/current error observed (as far as we know) in most r-z based cylindrical particle-in-cell codes. We then extend ideas included in the SMILEI documentation to find an improved method for computing the on-axis cylindrical field components for the first-order azimuthal harmonic; this change is shown to eliminate the previously observed spurious field growth. Finally, we discuss an improved implementation of absorbing boundary conditions in the quasi-3D geometry.

2.2.1 On-axis charge modification for uniform radial charge and current profiles

In cylindrical particle-in-cell codes, weighting of charges near the axis requires special treatment. If charges are loaded uniformly over the entire simulation domain, a spike in the charge (and current) density will be present in the first few cells near the axis due to the folding in of charge and current deposited in the negative r guard cells.[2] Thus, we seek to know if there is a proper initial charge normalization for particles near the axis such that the initial charge density is uniform throughout the simulation and charge is conserved, i.e., Gauss's law is maintained for the duration of the simulation. This work is extended from an idea proposed by Dan F. Gordon in a private communication.

First, it is important that we understand the discretization of fields and particles in OSIRIS for different geometries and interpolation levels. The particle position for each spatial dimension in OSIRIS is referenced from the center of what we refer to as a particle cell. The reference cell is stored as an integer (`ix`) that represents how many cells you are away from the left/lower/bottom boundary, and the offset from the cell center is stored as a floating point number ($-0.5 \leq \mathtt{x} < 0.5$). Storing the position as two separate numbers permits greater accuracy of the particle position for cells whose corner are located at positions with large values. In the current version of OSIRIS, the boundary of a particle cell is always fixed to the location of a simulation boundary; this is true regardless of the interpolation order. This simplifies and improves the efficiency of the communication pattern for particles.

The location of where field quantities in each direction are defined in relation to a reference particle cell varies depending on the interpolation level. In addition, the number of guard cells needed to push particles depends on the interpolation order. The width of a particle

[2]Interestingly, this is not an issue for particles initialized away from the axis that drift towards the cylindrical axis. Though we do not yet understand the exact reason for this, it appears to be a product of the charge-conserving current deposit scheme. In addition, a thermal plasma with immobile ions develops no spurious radial electric field near the axis over time. However, when utilizing the charge initialization scheme described in this section, a negative radial electric field (repulses electrons) proportional to r^{-1} develops for a thermal plasma with immobile ions as electrons initialized near the axis eventually drift away (and are replaced); this is an area for future work.

is $l + 1$, where l is the particle order: one for linear particles, two for quadratic particles, etc. Thus, odd (even) order particle shapes are an even (odd) number of cells wide in each direction. A field cell within the Yee lattice in OSIRIS is defined with ρ (charge) located at the lower/left/bottom corner, and E_x is displaced a half-cell forward in the x-direction (the E_y and E_z fields are also displaced by a half-cell in the y and z directions, respectively, in two and three dimensions). In OSIRIS the location of the field cells are aligned with the particle cells for odd particle orders and shifted by a half-cell for odd particle orders. This adjustment permits the use of similar code for field interpolation onto particles for any interpolation order.

However, in cylindrical geometry, the algorithm for the Maxwell solver assumes that the E_r field component is always defined at the cylindrical axis, or equivalently that ρ of the first cell be defined at $r = -\Delta r/2$. For this reason, the global simulation boundaries (which align with particle reference cells) are shifted in the r-direction from specified values to accommodate this. For example, for an input file requesting boundaries in r of $[0, r_{\max}]$, the actual boundaries will be $[-\Delta r/2, r_{\max} - \Delta r/2]$ and $[-\Delta r, r_{\max} - \Delta r]$ for odd and even interpolation levels, respectively. Note that this means that the actual simulation domain is $\Delta r/2$ or Δr smaller than $r_{\max}$ for odd and even interpolation orders, respectively.

Figure 2.8 shows the simulation discretization for two-dimensional runs in Cartesian and cylindrical geometries for odd and even interpolation levels. In this example the simulation domain extends from 0 to $2\Delta x$ (or the appropriate Δy, Δz or Δr) with 2 cells in each direction. Solid black lines represent cells to which the particles are indexed, with the center of the cell—where dotted lines intersect—corresponding to an OSIRIS particle position of $\mathtt{x} = (0.0, 0.0)$ and cell index $\mathtt{ix} = (1, 1)$. Therefore, the single black dot in Fig. 2.8 for each scheme represents a hypothetical particle located in particle cell $\mathtt{ix} = (1, 1)$ and with offset from center $\mathtt{x} = (0.25, 0.25)$. The green lines indicate the most lower-left field cell, for which red field quantities are also shown.

In Cartesian coordinates, we see that the particle cell boundaries in Fig. 2.8(a) and (b)

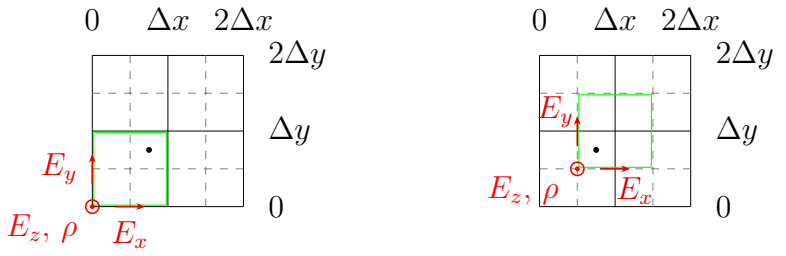

(a) Cartesian, odd

(b) Cartesian, even

(c) Cylindrical, odd

(d) Cylindrical, even

Figure 2.8: Field and particle discretization in OSIRIS for different geometries and particle interpolation levels. Black lines are cells about the center of which particles are indexed, with intersecting dotted lines indicating the center of those cells. Red field quantities are shown that correspond to the first field cell, which is outlined in green. A single black dot represents a hypothetical particle indexed as $\texttt{ix} = (1, 1)$ and $\texttt{x} = (0.25, 0.25)$.

are aligned with the requested simulation space boundaries, but that the physical position of the fields are shifted by a half-cell from particle cells for even interpolation order. However, in cylindrical coordinates the global simulation boundaries in the r-direction are shifted either by $-\Delta r/2$ or $-\Delta r$ from the requested domain of $0 \leq r \leq 2\Delta r$. This is done because the OSIRIS implementation always aligns particle cell boundaries with simulation boundaries, but the quasi-3D algorithm requires that E_r in the lower-most field cell be spatially defined at $r = 0$. With this adjustment, we can see that E_r is indeed located at $r = 0$ for both interpolation levels in Fig. 2.8(c) and (d).

Next, we examine how to correctly initialize the charge weighting for particles near the axis. Regardless of geometry or interpolation order, particles within a given cell are initially distributed about (and referenced from) the cell center (i.e., where dotted lines cross). In this way, particles are never initialized on particle cell boundaries; this choice is mostly for pure convention, but also avoids any ambiguity for particles initialized on a global simulation boundary. For even interpolation order in the cylindrical geometry with N particles per cell in r, the first particle nearest the axis will be loaded at a distance of $\Delta r/2N$ from the axis. However, for odd interpolation order, where the particle cell is shifted by half a cell, the particle nearest the axis will either be at a distance of $\Delta r/N$ or $\Delta r/2N$ from the axis for an odd or even value for the number of particles per cell, N. This means that for odd interpolation order we have different locations of the particles for even/odd N. However, from here on we will work as if the first particle was located at $\Delta r/2N$.

Assume that particles are distributed at positions $r_i = (i-1/2)\Delta r/N$, where $i \in 1, 2, 3, \ldots$ is the particle index. Let the n^{th} charge site, ρ_n, located above positive r be defined at position $R_n = (n - 1/2)\Delta r$, where $n \in 1, 2, 3, \ldots$. [In Fig. 2.8(c) and (d), the lower-most charge value (field cell index 1) is located at negative r, so the n used here actually differs by one from the OSIRIS cell index.] We then seek a function $Q(r_i)$ such that the total charge density in any cell is ρ for a uniform initialization of particles, keeping in mind that any charge deposited into a guard cell with $n \leq 0$ is added ("folded") back into the main simulation

space. The standard charge given to particles in OSIRIS, assuming one particle per cell in z and ϕ, is

$$Q_{\text{std}}(r_i) = \frac{\rho}{N} r_i. \tag{2.1}$$

The need for the dependence on r_i for constant density is clear for cells outside the first few cells as the size of particles scales with $2\pi r_i \Delta r$.

However, for the first cell this dependence needs to be altered as charge is folded back. We instead assume the charge for particles contained in just the first ℓ half-cells, for interpolation order ℓ, has an r-dependence whose order matches the interpolation level:

$$Q(r_i) = \begin{cases} \dfrac{\rho}{N} \displaystyle\sum_{j=1}^{\ell} \alpha_j \left(r_i\right)^j & r_i < \dfrac{\ell \Delta r}{2} \\ \dfrac{\rho}{N} r_i & r_i \geq \dfrac{\ell \Delta r}{2} \end{cases}. \tag{2.2}$$

The coefficients α_j can be solved for by equating the charge density of $\ell + 1$ cells, e.g., $\rho_1 = \rho_2$ and $\rho_2 = \rho_3$ for quadratic interpolation. Symbolic math software is especially helpful in solving for these coefficients. For the discretization used in OSIRIS, solving these equations yields

$$\alpha_1 = \begin{cases} \dfrac{2}{3} - \dfrac{2}{3N^2} & \text{even } N \\ \dfrac{2}{3} & \text{odd } N \end{cases} \tag{2.3}$$

for linear interpolation with $N \geq 2$. For quadratic interpolation with $N \geq 2$ we have

$$\begin{aligned} \alpha_1 &= \frac{7 - 22N^2}{24N^2(N^2-1)}, \\ \alpha_2 &= \frac{5(4N^2-1)}{4(4N^4 - 5N^2 + 1)}. \end{aligned} \tag{2.4}$$

The expressions for cubic interpolation once again depend on the parity of the number of

particles per cell and are given as

$$\alpha_1 = \begin{cases} \big[8680 - 294460N^2 - \\ \quad 6N^4(450036N^8 - 895320N^6 + 367533N^4 + 134135N^2 - 107514)\big] / \\ \quad \big[27N^2(3N-2)(3N+2)\cdot \\ \quad (9558N^8 - 26199N^6 + 15225N^4 - 2236N^2 + 3232)\big] & \text{even } N \\ \dfrac{-300024N^8 + 263520N^6 + 138738N^4 - 46310N^2 + 2996}{27(9558N^8 - 38043N^6 + 34601N^4 - 6457N^2 + 341)} & \text{odd } N \end{cases},$$

$$\alpha_2 = \begin{cases} \dfrac{2}{3}\dfrac{547074N^{10} - 1151361N^8 + 605937N^6 - 25518N^4 + 22114N^2 - 3496}{86022N^{10} - 274023N^8 + 241821N^6 - 81024N^4 + 38032N^2 - 12928} & \text{even } N \\ \dfrac{2}{9}\dfrac{182358N^8 - 346023N^6 + 169111N^4 - 11847N^2 - 339}{(9N^4 - 10N^2 + 1)(1062N^4 - 3047N^2 + 341)} & \text{odd } N \end{cases},$$

$$\alpha_3 = \begin{cases} -\dfrac{56}{27}\dfrac{80676N^{10} - 137376N^8 + 47673N^6 + 5997N^4 + 2926N^2 - 196}{86022N^{10} - 274023N^8 + 241821N^6 - 81024N^4 + 38032N^2 - 12928} & \text{even } N \\ -\dfrac{56N^2(8964N^6 - 13536N^4 + 4601N^2 + 301)}{27(9N^4 - 10N^1 + 1)(1062N^4 - 3047N^2 + 341)} & \text{odd } N \end{cases}. \tag{2.5}$$

Quartic interpolation is not yet implemented for quasi-3D Osiris.

The above method, however, is not valid for $N = 1$. In this case, all charges out to the radial boundary must actually be modified (rather than just those in the first ℓ half-cells), such that

$$Q(r_i) = \rho\beta_i r_i \tag{2.6}$$

for all values of i. However, when solving for the appropriate β_i coefficients, they converge to unity very quickly with increasing i. For linear interpolation with $N = 1$, there is actually no density spike on axis and all cells have equal charge ρ. For quadratic interpolation with $N = 1$ we let the first four coefficients vary from unity (i.e., $\beta_i = 1$ for $i \geq 5$) and solve those four coefficients by ensuring equal charge in the first five cells.[3] We then obtain the coefficients

$$\beta = \left\{ \frac{985}{1393}, \frac{607}{597}, \frac{6953}{6965}, \frac{9753}{9751}, 1, 1, \ldots \right\} \tag{2.7}$$

[3]Another way to generate these first four coefficients would be to let β_i vary for all values of i in the simulation (i.e., out to $r_{\max}$), solve for all the β_i coefficients by ensuring equal charge in all cells, then set $\beta_i = 1$ for $i \geq 5$. Solving for the coefficients in this way, we observed slight charge error in two cells, whereas the method described in the body text exhibited error mainly in one cell.

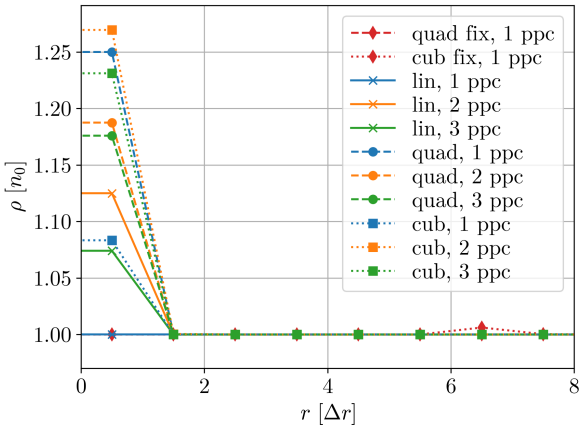

Figure 2.9: Deviations from uniform density without charge modification near the axis in cylindrical geometry for various interpolation orders and particles per cell. The first two lines show the only cases not entirely corrected with the new scheme.

for quadratic interpolation. Doing the same for cubic interpolation but with six variable coefficients and ensuring equal charge in seven cells, we find that

$$\beta = \left\{ \frac{10121}{11087}, \frac{11592}{11087}, \frac{32249}{33261}, \frac{11340}{11087}, \frac{10885}{11087}, \frac{11248}{11087}, 1, 1, \ldots \right\}. \tag{2.8}$$

Although the coefficients were derived under an assumption of uniform density, the idea can be applied for a non-uniform density. However, since a properly resolved density profile should not vary much in density from one cell to another we assume that within the first cell the density is constant. We thus perform these same calculations for any initialization, regardless of the density profile.

Once again, the charge correction needs to only be applied when initializing particles near the axis, i.e., at time step zero or at the front of a moving window. The improvement in the on-axis charge density is shown in Fig. 2.9 for various interpolation orders and particles per cell. Note that only for linear interpolation with one particle per cell was the density originally correct. The first two lines show quadratic and cubic interpolation, respectively, in the new scheme with one particle per cell. Only for these two cases is the density not completely

uniform, due to the truncation of the coefficients. For quadratic (cubic) interpolation there is a 0.002% (0.6%) increase in the charge density in the fifth (seventh) cell, one cell after the coefficients are truncated. The error is only visible for the cubic case, where we use no more than six cells due to rapidly diminishing returns in the sequence convergence (and to avoid using more coefficients than the minimum size of the simulation space). As such, it is advised to use more than one particle per cell in r when running with cubic interpolation.

In addition to the charge spike observed at the axis without the proper initialization, we also expect the on-axis fields and currents to exhibit errors. In Fig. 2.10 we show simulation quantities for an electron beam with a flat-top density profile free streaming forward at relativistic speed. We expect the radial electric field to be equal to $\rho r/2$ for a beam with charge density ρ, and in (a) we show the error in the observed electric field with and without the correction. Without the correction, we observe field artifacts corresponding to extra negative charge being located at the axis [see uncorrected charge density in Fig. 2.9]. We see a similar spike in the on-axis forward current in (c) without the correction. The (b) radial electric field and (d) charge density are also shown over the entire simulation domain for the corrected case, which are as to be expected.

2.2.2 Satisfying Gauss's law at the axis

Charge is conserved implicitly in OSIRIS; we do not explicitly solve Gauss's law, but it is implied due to the charge-conserving current deposit. The location of the fields, charge density, and current density on the grid are defined so that Gauss's law and the continuity equation are second-order accurate with cell size. In this configuration, Gauss's law is centered on points where the charge ρ is defined. In Cartesian coordinates, if $\nabla \cdot \boldsymbol{E} - \rho = 0$ is satisfied at discrete points across the entire simulation (i.e., where ρ is defined), then it is also satisfied at any point (i.e., in-between where ρ is defined) as long as quantities are assumed to vary linearly between grid points. This is because the divergence operator is linear in Cartesian coordinates, and we can take any linear combination of Gauss's law to be true. The same is

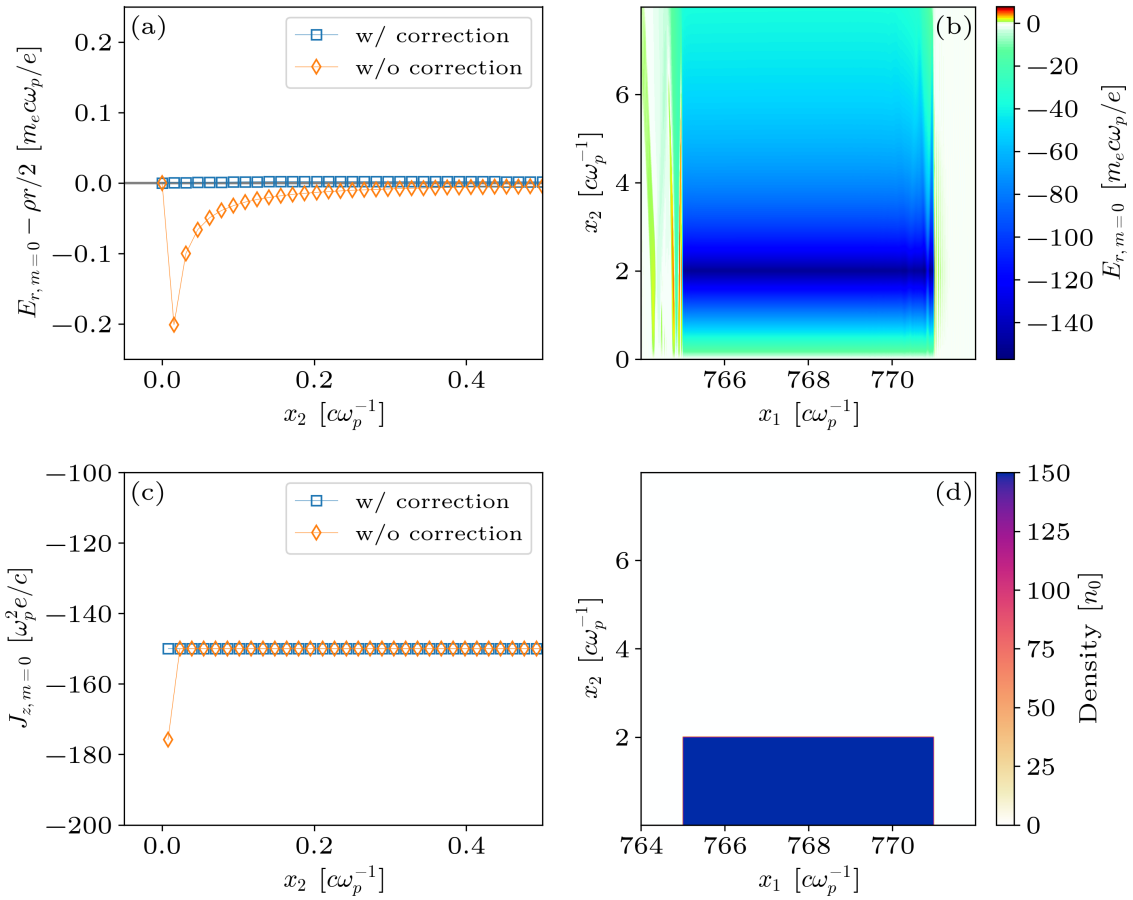

Figure 2.10: Simulation quantities for an electron beam free streaming relativistically in the positive x_1-direction. Lineouts along the beam center of (a) the error in $E_{r,m=0}$ and (c) observed $J_{z,m=0}$ with and without the charge correction. View of entire simulation domain for the corrected case showing (b) $E_{r,m=0}$ and (d) charge density.

true of the continuity equation.

However, this is not the case for cylindrical coordinates due to the nonlinearity of the divergence operator. Thus, the OSIRIS algorithm only ensures that Gauss's law and the continuity equation are satisfied at the grid points where ρ is defined. For this discussion we will refer to scalar (ρ), Cartesian (E_z, B_z) and cylindrical (E_r, E_ϕ, B_r, B_ϕ) fields. First we require a relation describing the field quantities for various modes right at the axis. As mentioned in Ref. [125], any scalar or Cartesian field quantity is single-valued at the axis, resulting in the necessary condition

$$\left.\frac{\partial S}{\partial \phi}\right|_{r=0} = 0, \quad \forall \phi \tag{2.9}$$

for any given field quantity S. Since quantities are expanded for mode m as $\Re\left\{S^m(r,z)e^{-im\phi}\right\}$, this means that the only nonzero mode at the axis for scalar and Cartesian quantities is mode 0. However, it turns out that mode 1 is the only nonzero component at the axis for cylindrical fields.

While rest of our derivations here apply to $\boldsymbol{E}$, $\boldsymbol{B}$ and $\boldsymbol{J}$, we will work mainly with the electric field. The analysis below benefited greatly from the documentation for the SMILEI particle-in-cell code [106]. We can express the transverse electric field as

$$\boldsymbol{E}_\perp = E_x\boldsymbol{e}_x + E_y\boldsymbol{e}_y = (E_r\cos\phi - E_\phi\sin\phi)\boldsymbol{e}_x + (E_r\sin\phi + E_\phi\cos\phi)\boldsymbol{e}_y. \tag{2.10}$$

Then evaluating Eq. (2.9) for E_x, we obtain

$$\cos\phi\left(\frac{\partial E_r}{\partial\phi} - E_\phi\right) - \sin\phi\left(\frac{\partial E_\phi}{\partial\phi} + E_r\right) = 0, \quad \forall\phi, \tag{2.11}$$

or equivalently,

$$\begin{aligned} \frac{\partial E_r}{\partial\phi} - E_\phi &= 0, \\ \frac{\partial E_\phi}{\partial\phi} + E_r &= 0. \end{aligned} \tag{2.12}$$

We see that for mode 0, Eq. (2.12) yields $E_r^0 = E_\phi^0 = 0$. However, evaluating for a given mode $m \geq 1$ gives

$$\begin{aligned} E_r^m &= \frac{iE_\phi^m}{m}, \\ E_r^m &= imE_\phi^m. \end{aligned} \tag{2.13}$$

The above equations can only be simultaneously satisfied for $m = 1$, meaning that while $E_r^m = E_\phi^m = 0$ for $m \geq 2$, we must have

$$E_r^1 = iE_\phi^1 \tag{2.14}$$

at the axis, which implies

$$\begin{aligned} E_{r,\mathrm{Re}}^1 &= -E_{\phi,\mathrm{Im}}^1, \\ E_{r,\mathrm{Im}}^1 &= E_{\phi,\mathrm{Re}}^1. \end{aligned} \tag{2.15}$$

Recall, the above Eq. (2.14) must also be true for the current $\boldsymbol{J}$ and magnetic field $\boldsymbol{B}$.

We also require that Gauss's Law holds (along with the continuity equation and divergence of $\boldsymbol{B}$). We will show that the E_r term at $r = 0$ is excluded by the divergence operator centered at $r = \frac{dr}{2}$. The divergence of a purely radial field is

$$\nabla \cdot E_r \boldsymbol{e}_r = \frac{1}{r}\frac{\partial}{\partial r}(rE_r) = \frac{E_r}{r} + \frac{\partial E_r}{\partial r}. \tag{2.16}$$

Since the charge in our scheme is centered at radial positions $r = \frac{dr}{2}$, $\frac{3dr}{2}$, ..., let us consider Eq. (2.16) at $r = \frac{dr}{2}$. We evaluate both versions of the derivative and show they are equivalent:

$$\nabla \cdot E_r \boldsymbol{e}_r = \frac{2}{dr}\left(\frac{dr \cdot E_r(r = dr) - 0 \cdot E_r(r = 0)}{dr}\right) = \frac{2E_r(r = dr)}{dr}, \tag{2.17}$$

or

$$\nabla \cdot E_r \boldsymbol{e}_r = \frac{E_r(r = dr) + E_r(r = 0)}{2} \cdot \frac{2}{dr} + \frac{E_r(r = dr) - E_r(r = 0)}{dr} = \frac{2E_r(r = dr)}{dr}. \tag{2.18}$$

In either case, the $E_r(r = 0)$ field does not contribute to Gauss's law centered at $r = \frac{dr}{2}$.

We have verified that the quasi-3D OSIRIS code is in fact charge conserving at each of the simulation points where ρ is defined. However, we now see that $E_r(r=0)$ could take on any value and this would still be true. In particular, although a formula is derived for calculating $E_r^1(r=0)$ in Ref. [30] (remember, $E_r^m(r=0)=0$ for $m\neq 1$), growth (sometimes large) of the $E_r^1(r=0)$ field is observed in nearly every quasi-3D simulation. In some simulations, after many time steps the on-axis field is an order of magnitude larger than the field one cell away. This results in erratic particle motion. It is largely in an effort to mitigate this growth that we pursue an alternate calculation of the on-axis fields.

Although Gauss's law is only strictly satisfied at grid points where ρ is defined, i.e., at the position $r=\frac{dr}{2}$ for our discretization, we now endeavor to ensure that it is also satisfied at $r=0$ (along with the continuity equation and divergence-free $\boldsymbol{B}$). For the $m=1$ mode, we know that $\rho^1(r=0)=0$, so we require $\nabla\cdot\boldsymbol{E}^1=0$. The z-component is also zero at the axis, so we are left with

$$\nabla\cdot\boldsymbol{E}^1=\frac{1}{r}\frac{\partial}{\partial r}(rE_r^1)-\frac{iE_\phi^1}{r}=\frac{E_r^1-iE_\phi^1}{r}+\frac{\partial E_r^1}{\partial r}=0. \tag{2.19}$$

Assuming that Eq. (2.14) is true, the divergence operator at the axis then reduces to

$$\nabla\cdot\boldsymbol{E}^1=\left.\frac{\partial E_r^1}{\partial r}\right|_{r=0}=0. \tag{2.20}$$

Although $\boldsymbol{E}$ and $\boldsymbol{J}$ are defined at different coordinates than is $\boldsymbol{B}$, Eq. (2.20) can be used to find either the r- or ϕ-component of a cylindrical field directly at the axis through Eq. (2.14).

Next, we state how Eq. (2.20) is enforced in the code for the two different field discretizations. First, we work with the electric field, where E_r is located on the axis. We can build a second-order Taylor expansion for the known field quantities at $r=dr$ and $2dr$ using the unknown field quantity and derivatives at the axis. If we define $u(r)=E_r^1(r)$, then we can write

$$\begin{aligned} u(dr)&\approx u(0)+dr\,u'(0)+\frac{dr^2}{2}\,u''(0),\\ u(2dr)&\approx u(0)+2dr\,u'(0)+2dr^2\,u''(0). \end{aligned} \tag{2.21}$$

Since we want to enforce $u'(0) = 0$, we can eliminate $u''(0)$ from the above equations and solve for $u'(0)$ as

$$u'(0) = \frac{-3\,u(0) + 4\,u(dr) - u(2dr)}{2dr} = 0. \tag{2.22}$$

Solving for the quantity at the axis yields

$$u(0) = \frac{4}{3}\,u(dr) - \frac{1}{3}\,u(2dr). \tag{2.23}$$

The same approach can be taken using the points $r = \frac{dr}{2}$ and $\frac{3dr}{2}$ to obtain

$$u(0) = \frac{9}{8}\,u\left(\frac{dr}{2}\right) - \frac{1}{8}\,u\left(\frac{3dr}{2}\right). \tag{2.24}$$

We use Eq. (2.23) to give $E_r^1(r = 0)$. Then once $B_r^1(r = 0)$ is determined via Eq. (2.24), we can use Eq. (2.14) to calculate $B_\phi^1(r = 0) = -iB_r^1(r = 0)$.

Now that the quantities E_r^1 and B_ϕ^1 are known at the axis, we want to ensure that when electric fields are interpolated onto particles in the pusher, the relation in Eq. (2.14) will hold. We have $E_r^1(r = 0)$ specified by Eq. (2.23), and $E_\phi^1\left(r = \frac{dr}{2}\right)$ is dictated by the normal field solver (modifying it will void charge conservation). However, we can modify E_ϕ^1 in the guard cells at negative r values to satisfy the relation. We simply want the average of the two E_ϕ^1 values to be proportional to E_r^1, or

$$\frac{E_\phi^1\left(-\frac{dr}{2}\right) + E_\phi^1\left(\frac{dr}{2}\right)}{2} = -iE_r^1(0), \tag{2.25}$$

or

$$E_\phi^1\left(-\frac{dr}{2}\right) = -2iE_r^1(0) - E_\phi^1\left(\frac{dr}{2}\right). \tag{2.26}$$

This is also done for the guard cell at $r = -\frac{3dr}{2}$.

After implementing these changes into OSIRIS, we saw dramatic improvement in the on-axis E_r^1 field. Primarily, an artifact that was present in many simulations was that plasma would be expelled from the axis near the very leading edge of the laser fields. This behavior, seen in Fig. 2.11(a), is absent when using the on-axis field correction described above, as

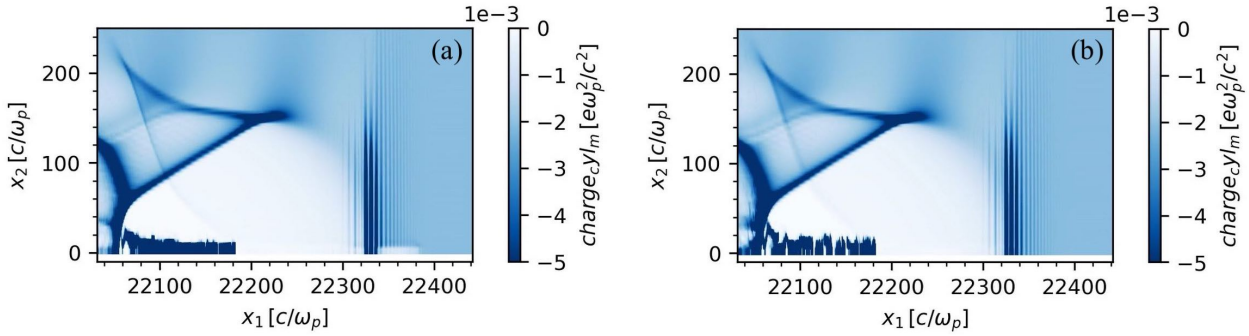

Figure 2.11: Laser wakefield simulation (a) without and (b) with the on-axis field corrections described in Sec. 2.2.2. The erroneous expulsion of plasma from the axis before the laser pulse is removed with the fix.

seen in Fig. 2.11(b). (Note that this on-axis behavior is absent for both cases if no laser is present.)

In addition, the single particle motion in a laser improved with this correction. In the current state of the code, when test charges propagate inside a laser, they appear to move as expected near the axis. However, when their charge is non-negligible, we observe a slight attraction toward the axis over many laser cycles. We thought of somehow modifying the J^1_ϕ current in the guard cells at negative r (previous to folding the current back into the main simulation space) akin to that done for the electric field in Eq. (2.26), but any modification to the J^1_ϕ values makes the simulation not charge conserving. Future work involves determining the cause of the slight on-axis attraction for particles moving in a laser field. It is unclear whether the issues of on-axis field growth and particle attraction are a result of the particular field discretization used in OSIRIS, and one option would be to test the discretization used in Lifschitz [125] by shifting the field quantities a half-cell in r. Another possible cause is that since Gauss's law is only satisfied where ρ is defined (and now at $r = 0$), this attraction is due to small divergence errors between those points, errors that are only noticeable in sensitive systems over long duration. The currents and fields acting on particles are computed using some interpolation scheme, which assumes that these quantities vary accordingly in

the space between the grid points. Thus, errors in Gauss's law are likely present throughout the simulation space away from where ρ is defined, but it is possible that this error is only significant close to the axis.

2.2.3 Absorbing boundary conditions in quasi-3D

Various absorbing boundary conditions (ABCs) have been devised to behave as "open" boundaries, where outgoing radiation exists as if the simulated region was a vacuum with infinite extent beyond the boundary. The simplest versions of these impose a condition on the first cell outside of the simulation domain to absorb outgoing waves. One of the most notable and straightforward of these ABCs was proposed by Mur [128], where he factored the wave equation into two operators corresponding to waves traveling in opposite directions. Taylor expanding about normal incidence and setting the appropriate operator to zero at the correct boundary will absorb waves exiting that region, with diminishing returns for increased angle of incidence. The Lindman ABC [38] uses a more sophisticated expansion in comparison to Mur and in turn attains greater absorption, especially for evanescent waves [129]. Here we outline the Mur ABC implementation in cylindrical coordinates under the modal expansion in ϕ.

First, we have that the wave equation in cylindrical coordinates for quantity f is

$$\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial f}{\partial r}\right)+\frac{1}{r^2}\frac{\partial^2 f}{\partial \phi^2}+\frac{\partial^2 z}{\partial z^2}-\frac{1}{c^2}\frac{\partial^2 f}{\partial t^2}=0, \tag{2.27}$$

where we have that $\frac{\partial^2 f}{\partial \phi^2}=-m^2 f$ for a given mode m. To formulate the Mur boundary condition in the z-direction, we factor Eq. (2.27) into the two operators

$$\begin{aligned} G^- &= D_z-\frac{D_t}{c}\sqrt{1-s^2}, \\ G^+ &= D_z+\frac{D_t}{c}\sqrt{1-s^2}, \end{aligned} \tag{2.28}$$

with

$$s=\frac{\sqrt{\frac{1}{r}D_r\left(rD_r\right)-\frac{m^2}{r^2}}}{D_t/c}, \tag{2.29}$$

where we have used the notation $D_q \equiv \frac{\partial}{\partial q}$. Factored in this way, Eq. (2.27) can be rewritten as $G^+ G^- f = 0$.

We can eliminate outgoing waves at the simulation boundaries by applying G^- and G^+ at the left and right boundaries, respectively; i.e., we require that

$$G^- f\big|_{z=0}, \quad G^+ f\big|_{z=h} \tag{2.30}$$

for a simulation with $z \in [0, h]$. In order to implement this condition, we must Taylor expand the square roots in Eq. (2.28) assuming small s (expanding about normal incidence). Keeping terms of first order in s and taking another derivative in time yields the operators

$$D_t G_1^{\pm} = D_{zt} \pm \frac{1}{c} D_t^2 \mp \frac{c}{2} \left[\frac{1}{r} D_r \left(r D_r \right) - \frac{m^2}{r^2} \right]. \tag{2.31}$$

The ABC is implemented by modifying the first cell outside the main simulation space to impose the condition

$$D_t G_1^{\pm} f = 0. \tag{2.32}$$

However, as in Mur's original formulation, we only need to solve Eq. (2.32) for electric field quantities tangential to the boundary, i.e., E_r and E_ϕ. Consider a simulation extending from 0 to h with N cells in the z-direction, with the lower-left corner of cell $(i, j) = (1, 1)$ located at $z = 0$ and $r = -\Delta r/2$ for radial step size Δr. The OSIRIS field solver computes $\boldsymbol{B}$ over the cell range $i \in [0, N+1]$ and computes $\boldsymbol{E}$ over the cell range $i \in [1, N+1]$. The field solve in the radial direction is standard except for the cell with $j = 1$, where special axial boundary conditions must be imposed. We then impose Eq. (2.32) at $i = 0$ and $i = N + 2$. Since this involves second derivatives in t and r and first derivatives in z, we solve for $f_{0,j}^{n+1}$ and $f_{N+2,j}^{n+1}$ using derivatives centered about $f_{1/2,j}^{n}$ and $f_{N+3/2,j}^{n}$, respectively, where $f_{i,j}^{n}$ is the quantity f located in cell (i, j) and at the n^{th} time step. Taking finite-difference representations of the derivatives in Eq. (2.31) and solving for the appropriate quantity in Eq. (2.32) yields the

left-boundary value as

$$
\begin{aligned}
f_{0,j}^{n+1} = \frac{1}{2\Delta r^2 r_j(\Delta z + \Delta t)} \Big\{ & \Delta z \Delta t^2 \left[r_{j-1/2} \left(f_{0,j-1}^n + f_{1,j-1}^n \right) + r_{j+1/2} \left(f_{0,j+1}^n + f_{1,j+1}^n \right) \right] \\
& + \left[2\Delta z\, r_j \left(2\Delta r^2 - \Delta t^2 \right) - \frac{\Delta r^2 \Delta t^2 \Delta z}{r_j} m^2 \right] \left(f_{0,j}^n + f_{1,j}^n \right) \\
& - 2\Delta r^2 \Delta z\, r_j \left(f_{0,j}^{n-1} + f_{1,j}^{n-1} + f_{1,j}^{n+1} \right) \\
& + 2\Delta r^2 \Delta t\, r_j \left(f_{0,j}^{n-1} - f_{1,j}^{n-1} + f_{1,j}^{n+1} \right) \Big\},
\end{aligned}
\tag{2.33}
$$

and the right-boundary value is given by

$$
\begin{aligned}
f_{N+2,j}^{n+1} = \frac{1}{2\Delta r^2 r_j(\Delta z + \Delta t)} \Big\{ & \Delta z \Delta t^2 \big[r_{j-1/2} \left(f_{N+1,j-1}^n + f_{N+2,j-1}^n \right) \\
& \qquad + r_{j+1/2} \left(f_{N+1,j+1}^n + f_{N+2,j+1}^n \right) \big] \\
& + \left[2\Delta z\, r_j \left(2\Delta r^2 - \Delta t^2 \right) - \frac{\Delta r^2 \Delta t^2 \Delta z}{r_j} m^2 \right] \left(f_{N+1,j}^n + f_{N+2,j}^n \right) \\
& - 2\Delta r^2 \Delta z\, r_j \left(f_{N+1,j}^{n-1} + f_{N+1,j}^{n+1} + f_{N+2,j}^{n-1} \right) \\
& + 2\Delta r^2 \Delta t\, r_j \left(-f_{N+1,j}^{n-1} + f_{N+1,j}^{n+1} + f_{N+2,j}^{n-1} \right) \Big\},
\end{aligned}
\tag{2.34}
$$

where m is the cylindrical mode number and r_j is the radial distance of quantity $f_{*,j}$ (will differ between E_r and E_ϕ).

The process of enforcing the Mur boundary condition is as follows for the left boundary, beginning first with the $\boldsymbol{B}$ field solve:

1. Set all $\boldsymbol{B}$-field components to the left of (inclusive) $i = -1$ to zero
2. Solve for E_z at $i = 0$ in the standard fashion
3. Solve for E_r and E_ϕ at $i = 0$ using Eq. (2.33)
4. Impose the axial boundary condition for $\boldsymbol{E}$ at cell $(i, j) = (0, 1)$
5. Copy requisite E_r and E_ϕ field data into buffers

6. Set all $\boldsymbol{E}$-field components to the left of (inclusive) $i = -1$ to zero

Similarly, we give the process at the right boundary:

1. Solve for B_z at $i = N + 2$ in the standard fashion, also imposing the axial boundary condition for this component for cell $j = 1$

2. Set to zero B_r and B_ϕ at $i = N + 2$ and all $\boldsymbol{B}$-field components to the right of (inclusive) $i = N + 3$

3. Solve for E_r and E_ϕ at $i = N + 2$ using Eq. (2.34), also imposing the axial boundary condition for these components for cell $j = 1$

4. Copy requisite E_r and E_ϕ field data into buffers

5. Set to zero E_z at $i = N + 2$ and all $\boldsymbol{E}$-field components to the right of (inclusive) $i = N + 3$

When implemented as described above, normally incident outgoing waves are absorbed well by the boundary everywhere except for near the axis. In one test, a reflected wave was generated on the order of $\approx 10\%$ of the original amplitude that propagated back into the simulation in both the z and r directions for a laser with mode $m = 1$. This reflected wave disappeared when we set $m = 0$ in Eqs. (2.33) and (2.34) for all modes. This is because our Taylor expansion of Eq. (2.28) assumes small s. However, near the origin the $\frac{m^2}{r^2}$ term in Eq. (2.29) is very large. Thus the current implementation of the Mur boundary condition in OSIRIS follows Eqs. (2.33) and (2.34) but without the term multiplied by m.

Note that the Mur boundary condition as described here is only valid at boundaries in z. Other boundary conditions have been formulated for use at the outer radial boundary [130, 131], but these are not discussed in this dissertation.

2.2.4 Summary

The quasi-3D algorithm [30, 125] is a powerful tool that enables simulations with three-dimensional geometric effects at the computational cost of a two-dimensional simulation. The separation of modes also allows for some unique analysis techniques that are otherwise difficult in 3D Cartesian (see Sec. 3.1). However, the nonlinearity of spatial derivative operators and the handling of the central axis present several nuanced challenges. In this section we presented a charge initialization scheme that correctly yields the desired density at the axis, a modification to the field solver that eliminates spurious on-axis field growth by satisfying Gauss's law at the axis, and an implementation of the Mur boundary condition at the left and right simulation edges in z to absorb outgoing electromagnetic waves. These fixes are crucial developments to ensure stable, fully self-consistent simulations of laser–plasma interactions and are utilized in much of the work presented in the following chapter.

CHAPTER 3

Role of direct laser acceleration in underdense laser–plasma interactions

Traditional short-pulse laser wakefield acceleration (LWFA) has been shown to be effective in generating energetic beams of electrons [82, 9]. In LWFA, the radiation pressure of a short-pulse laser creates a plasma wave wake as it propagates over large distances through a tenuous plasma. The LWFA regime of interest is the nonlinear self-guided regime [67]. The wake is created by expelling essentially all of the plasma electrons, which then form a narrow sheath that surrounds what is referred to as an ion column. If the laser pulse is sufficiently short that it fully resides in the front half of the first wavelength (bucket) of the wakefield, a trailing beam of electrons placed at the rear of the first bucket will experience the accelerating and focusing phase of the wake (and no laser fields). However, if the laser duration is long enough such that the laser fields overlap the rear of the first bucket (or a later bucket), then the laser fields can affect the acceleration process.

If the laser fields fill just the first plasma bucket, many features of the short-pulse LWFA are preserved (e.g., wake excitation, self or external guiding and pump depletion), while the acceleration of the electrons at the back of the bucket is modified by the presence of the laser. The acceleration of electrons in the combined wakefields and laser fields can be complicated. For example, there is the possibility of betatron resonances. Inside the focusing fields of the ion column in a nonlinear wakefield, electrons execute periodic transverse motion at what is called the betatron frequency ($\omega_p/\sqrt{2\gamma}$). When the transverse oscillation ("bounce" or betatron) frequency of the electrons in the channel, ω_β, matches the Doppler-shifted laser

frequency, the electrons are accelerated by a process originally called direct laser acceleration (DLA) [78, 70, 79, 80, 1]. The resonance condition for DLA is thus $N\omega_\beta = \omega_0 - k_0 v_z$, where ω_0 and k_0 are the frequency and wavenumber of the laser, respectively, v_z is particle velocity in the laser propagation direction, and N is the harmonic number. This term has been expanded to also include the simultaneous acceleration in the total fields (longitudinal and transverse) inside the wake and any overlapping laser fields.

For a laser with a duration that overlaps many plasma buckets, the laser can self-modulate, i.e., split up into multiple beamlets one plasma period long. This leads to a plasma wave that can grow over multiple periods. While the concept of a wake relates to fields behind the driver, in what follows we also refer to the accelerating fields inside the laser as wakefields. This regime of laser-driven acceleration is called self-modulated laser wakefield acceleration (SM-LWFA) [55, 57, 132, 58, 59]. This process not only results in more trapped charge due to the increased number of plasma bubbles, but also results in enhanced betatron radiation from electrons experiencing DLA [133]. If the laser is sufficiently long and intense, an ion channel can form where the electrons are unable to return to the axis due to the sustained ponderomotive force of the laser [70, 134, 135]. The ions may also eventually begin to move, forming a near-hollow cylindrical channel at the back of the laser. As mentioned above, electrons moving in an ion channel (even partially evacuated) can gain significant energy from DLA alone in the absence of a longitudinal wakefield [133, 135].

DLA has been shown in many cases to be an important factor in high-energy electron generation in LWFA and SM-LWFA schemes [70, 68, 136, 1, 74, 135]. Many efforts have been made to more fully understand and enhance this process [137, 138, 139]. This chapter deals with more accurately quantifying and simulating the betatron resonance and DLA processes.

3.1 Accurate accounting of energy gain from direct laser acceleration

There is interest in understanding in detail the contribution from direct laser acceleration (DLA) to the energy gain of accelerated electrons in a laser wakefield accelerator (LWFA). Experimental signatures such as high-energy forking in electron spectra [140, 135, 74] are indicative of DLA, but we desire to quantitatively determine its impact via particle-in-cell simulation. The DLA process occurs when the transverse betatron oscillations of electrons in an ion channel are in resonance with the Doppler-shifted laser frequency. In the presence of a longitudinal wakefield, electrons are accelerated both in the forward direction by the wake and in the transverse direction by the laser field. The transverse energy gain from the laser field is converted into forward momentum by the $\boldsymbol{v} \times \boldsymbol{B}$ force of the laser.

In previous work [68, 1, 134, 135], the work done by the DLA (LWFA) process is assumed to be completely due to the transverse (longitudinal) electric field with respect to the propagation direction. For this discussion, the terms transverse and longitudinal do not mean where the divergence or curl vanishes. This assumption, although fairly accurate, does not take into account the fact that a finite-size laser pulse has a field component in the longitudinal direction and that the wakefield has components in the transverse direction. A Gaussian laser beam polarized in the x-direction and propagating in the z-direction has transverse and longitudinal field amplitudes of [68]

$$\begin{aligned} E_x &= E_0 e^{-(r/\sigma)^2} \sin(kz - \omega t) \\ E_z &\approx -\frac{2x}{k\sigma^2} E_0 \cos(kz - \omega t), \end{aligned} \tag{3.1}$$

where $r = \sqrt{x^2 + y^2}$, k and ω are the laser wavenumber and frequency, respectively, E_0 is the laser amplitude and σ is the beam waist. Note that the longitudinal laser field has a $\pi/2$ phase difference from the transverse component.

To determine the impact of this longitudinal field on the DLA of electrons, the ratio f of the average energy gained during a single betatron oscillation from the parallel ($\langle W_\parallel \rangle$) and

transverse ($\langle W_\perp \rangle$) fields of a Gaussian laser is given as [68]

$$f \equiv \frac{\langle W_\parallel \rangle}{\langle W_\perp \rangle} = -\frac{2c^2\sqrt{2\gamma}}{\omega\omega_p\sigma^2}, \tag{3.2}$$

where ω_p is the plasma frequency. We observe that the negative sign in Eq. (3.2) indicates that electrons in betatron resonance with a Gaussian laser pulse will always lose energy due to the longitudinal field component. This energy loss is neglected when calculating the DLA energy as only the work done by the perpendicular fields, which overestimates the DLA contribution in such simulations.

To more accurately compute the total energy gain (including longitudinal losses) from the DLA mechanism, we leverage the quasi-3D geometry of Osiris. In this geometry, cylindrical coordinates r, z and ϕ are used, with all field and current quantities split into modes based on their ϕ-dependence. When assuming a cylindrically symmetric laser spot size, all fields (both transverse and longitudinal) due to the wake are cylindrically symmetric and thus pertain to the $m = 0$ azimuthal mode. All fields (both transverse and longitudinal) for a linearly polarized laser with a symmetric spot pertain to the $m = 1$ azimuthal mode. Because of this distinction, the longitudinal fields corresponding to the wake and laser are separate quantities. Thus, instead of the traditional calculation of work done on an electron by DLA and LWFA,

$$\begin{aligned} W_{\text{LWFA}} = W_\parallel &= -e \int E_z \cdot v_z \, dt = -e \sum_m \int E_z^m \cdot v_z \, dt \\ W_{\text{DLA}} = W_\perp &= -e \int \boldsymbol{E}_\perp \cdot \boldsymbol{v}_\perp \, dt = -e \sum_m \int \boldsymbol{E}_\perp^m \cdot \boldsymbol{v}_\perp \, dt, \end{aligned} \tag{3.3}$$

we can more accurately calculate the total work as

$$\begin{aligned} W_{\text{LWFA}} &= -e \int \boldsymbol{E}^0 \cdot \boldsymbol{v} \, dt \\ W_{\text{DLA}} &= -e \int \boldsymbol{E}^1 \cdot \boldsymbol{v} \, dt, \end{aligned} \tag{3.4}$$

where a superscript on the electric field denotes the azimuthal mode.[1]

[1]Note that when we refer to the LWFA (or SM-LWFA) mechanism, we refer not only to the plasma

The work calculation as described by Eq. (3.4) was carried out in a similar manner using a quasi-static PIC code, QS-DLA, in Ref. [134], where it was found that the longitudinal laser field did indeed do negative work on electrons in betatron resonance for the parameters studied. The effect of a laser in a quasi-static approximation (QSA) code is implemented through the use of the ponderomotive guiding center (PGC) approximation [141]. In the PGC, high-frequency laser motion is averaged out, resulting in a ponderomotive force term that depends on gradients of the laser envelope. The laser is evolved through an envelope equation for the vector potential, where the index of refraction comes from the dynamics of plasma particles. To include DLA into a PGC-type framework, the envelope is multiplied by the high-frequency parts to get the vector potential, which is used to get the electric and magnetic fields. However, the QS-DLA code cannot capture self-injection of electrons (both in LWFA and SM-LWFA regimes), model pump depletion distances, or include non-PGC physics of interest. The simulations in Ref. [134] were performed only with an externally injected electron bunch.

The calculation in Eq. (3.4) can be performed using the fully self-consistent, relativistic PIC code Osiris for many additional cases of interest. In particular, the relative roles of LWFA and DLA were explored for an SM-LWFA in the absence of a trailing ion channel in Ref. [74], as is hereafter described.

---

wave wake left behind by the laser, but also to any plasma wave response located within the laser pulse. By the LWFA mechanism, we mean the energy gained due to fields from the plasma wave wake; these energy-contributing fields are predominantly longitudinal and azimuthally symmetric. The DLA mechanism refers to energy gained due to laser fields; such fields involve transverse and longitudinal fields oscillating at the laser frequency. For a short-pulse laser pulse that fills less than one bucket of the plasma wave wake, the DLA mechanism will be active only near the front of the first bucket, with the LWFA mechanism active in the plasma wave found behind the laser. However, for a longer pulse duration that exists over many plasma periods, both DLA and LWFA mechanisms may be present within the laser pulse.

3.1.1 Direct laser acceleration in a low-density self-modulated laser wakefield accelerator

Many high energy density science (HEDS) facilities such as the National Ignition Facility at the Lawrence Livermore National Laboratory, the OMEGA Laser at the University of Rochester, the Z-Machine at Sandia National Laboratories, and the Laser Mégajoule at the Commissariat à l'Energie Atomique are all coupled to picosecond-duration, kilojoule-class laser systems. Understanding the relationships between temperature, pressure, and density in the extreme environments created at these large facilities is crucial and has implications in planetary science [142], inertial confinement fusion [5], and laboratory astrophysics [143]. A limiting factor has been the quality of x-ray sources available to diagnose experiments. LWFA-driven sources have the potential to enable ultrafast resolution of dynamic experiments and provide measurements with unparalleled spatial resolution [133, 144, 145]. While the promise of such x-ray sources is evident, optimization of LWFA-driven x-rays with the picosecond-duration lasers co-located at HEDS facilities necessitates a detailed understanding of the underlying physics of electron beam generation mechanisms—not in the LWFA regime, but in the self-modulated LWFA (SM-LWFA) regime [58], which is as yet incomplete.

Prior work on SM-LWFA has attributed the copious charge of high-energy electrons observed in experiments to self-trapping and breaking of the longitudinal plasma wave [71, 132]. Electrons with energies larger than the dephasing-limited energy gain were also observed [146, 147], but were initially attributed to acceleration by the plasma wave. The role of DLA in LWFA was first suggested by Pukhov [78] and has been experimentally investigated only recently. In a quasi-blowout regime, the laser pulse was lengthened to overlap with a full plasma period, and electrons in the high-energy tail of the accelerated electron spectrum showed a fork-like splitting when dispersed perpendicular to the laser polarization direction [136, 1]. This fork-like structure was attributed to DLA based on PIC simulations, but the analysis did not include the contribution of the longitudinal field from the focused laser in the DLA process. In a high-density ($\sim 10^{20}$ cm^{-3}), short-pulse

(50 fs) SM-LWFA regime, the high-energy electron beam tail was experimentally attributed to DLA [148], but without any clear experimental signature. In a long-pulse (650 fs), high-intensity ($I = 3 \times 10^{20}$ W/cm^2) regime, DLA was inferred from PIC simulations to be the main acceleration mechanism in an ion channel [72]. The role of DLA in an SM-LWFA was anticipated in a recent experiment on developing a betatron-radiation-based x-ray source, but no direct experimental evidence for DLA was presented in that work [133, 73].

Here we discuss how the improvements described in Chapter 2 and later in this chapter to the quasi-3D algorithm of the particle-in-cell (PIC) code OSIRIS [28, 30] permit modeling of experiments performed at Lawrence Livermore National Laboratory. The experimental evidence together with the simulations indicate that DLA is the dominant contributor to the highest-energy electrons in the low-density ($\omega_p \ll \omega_0$) regime of SM-LWFA. In this regime, simulations show that the laser power—although it is slightly greater than that needed for relativistic self-focusing—is insufficient to produce a totally evacuated ion channel inside the laser pulse. The experimental work shows that, for a low plasma density ($\omega_0/\omega_p = 57$) and a 1-μm, nominally 1-ps laser with moderate amplitude $a_0 = 8.5 \times 10^{-8}\lambda I^{1/2} \approx 2.2$, the accelerated electrons exhibit a two-temperature distribution. Here a_0 is the normalized vector potential and λ [nm] is the laser wavelength. The full-scale quasi-3D OSIRIS simulations confirm that the longitudinal field of the plasma wave excited by the SM-LWFA process mainly contributes to the low-temperature portion of the spectrum, whereas DLA is the dominant acceleration mechanism for the high-energy (temperature) electrons. When the electrons are experimentally dispersed orthogonal to the laser polarization direction, a fork-like structure [1] characteristic of DLA is observed for electrons with energies above 60 MeV. This is the first direct experimental characterization, confirmed by quasi-3D PIC simulations, of DLA in an SM-LWFA in the picosecond, high-energy regime relevant to HEDS experiments.

The experiment was conducted on the Titan laser system at Lawrence Livermore National Laboratory (Fig. 3.1). Titan, a $0.7^{+0.3}_{-0.1}$ ps, 120 J, Nd:Glass laser, was focused with an $f/10$ off-axis parabolic mirror into a 10-mm, supersonic He gas jet with electron density

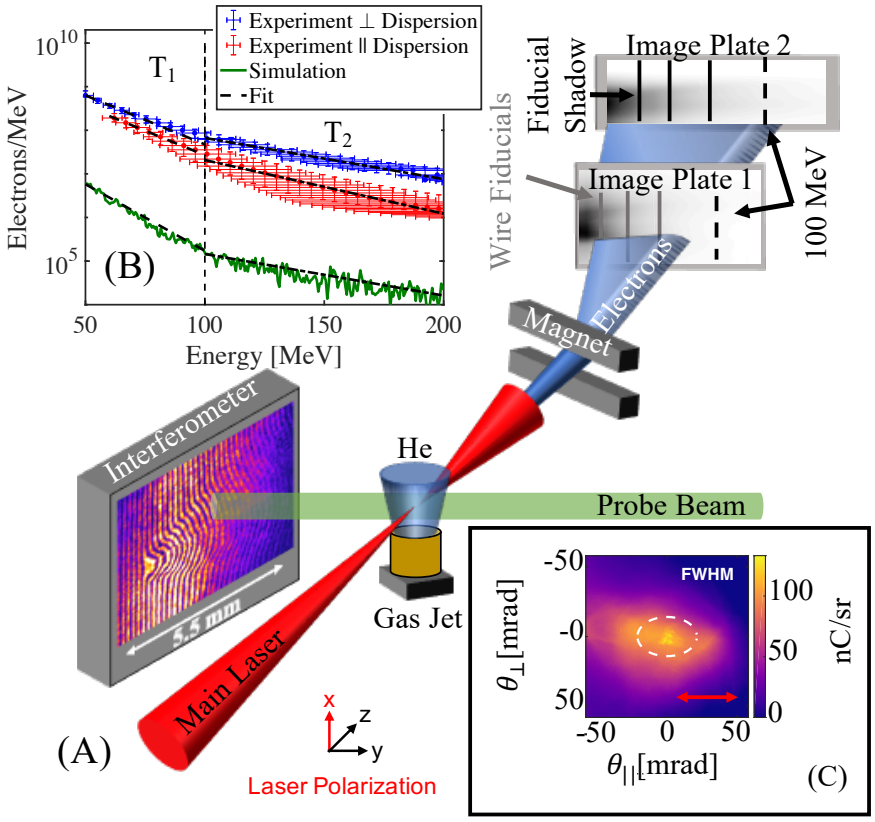

Figure 3.1: (a) Experimental setup. The electron beam is dispersed by a 0.6 T magnet onto two BAS-MS image plates after passing through three wire fiducials. A frequency-doubled probe beam is co-timed with the main pulse and provides on-shot interferometry of the plasma channel with a magnification of 3. (b) The electron energy spectrum for two different shots using identical laser and plasma parameters dispersed perpendicular and parallel to the laser polarization in blue and red, respectively, along with the simulated electron spectrum in green. All three spectra are fit to single-temperature distributions below (T_1) and above (T_2) 100 MeV; the two regions are separated by a dashed black line. The experimental spectra exhibit shot-to-shot variations of the single-shot laser system. (c) An undispersed electron beam profile.

$n_e = 3 \times 10^{17}$ cm^{-3} (measured using interferometry). This configuration created peak laser intensities reaching $I = 6.4 \times 10^{18}$ W/cm^2, in a spot with 50% of the total energy contained in a 30-μm radius. The ratio $P_{\text{peak}}/P_{\text{crit}} \approx 1.6$, where P_{peak} is the peak laser power and $P_{\text{crit}} = 17 \times 10^9 \left(\frac{\omega}{\omega_p}\right)^2 = 56$ TW is the critical power for relativistic self-focusing in the plasma [82, 67].

The electron beam is dispersed by a 0.6 T magnet with a large opening aperture (5 cm) perpendicular to the linear laser polarization. A smaller magnet (2.56 cm aperture) with a 1 T field is used to disperse the electrons parallel to the linear laser polarization. The dispersed electron signal passes through three wire fiducials—used to reduce error in the energy calculation [149]—and is captured on two BAS-MS image plates separated by 30 cm. The difference in the electron energy spectra in Fig. 3.1(b) is caused by shot-to-shot variation in the laser energy, pulse width, and quality of the high-power spot size. In an ideal case, these two spectra would be identical.

The non dispersed electron beam data from the experiments is shown in Fig. 3.2(a) which shows an elongation of the electron beam along the laser polarization axis. This trend is evident when comparing all shots using both the 4 and 10 mm gas jet diameter nozzles and ranging in plasma density from 3.5×10^{17} to 6.5×10^{18} cm^{-3}. The elongation along the linear laser polarization is evidence of the impact of DLA during the acceleration process as electrons which oscillate along the laser polarization will bunch up at the extremes of the oscillation [1].

Additionally, the electron energy spectra for several nozzle types and plasma densities is fit to a two-temperature spectrum similar to the one shown in Fig. 3.1. The results are presented in Fig. 3.2(b), where the measured hot and cold temperatures—T_1 and T_2, respectively—follow the similar trend of two distinct distributions seen above and below 100 MeV. The two-temperature distribution is hypothesized to be due to SM-LWFA acceleration mechanisms dominating the energy gain at lower final energies, whereas DLA begins to dominate the energy gain at around 60 MeV.

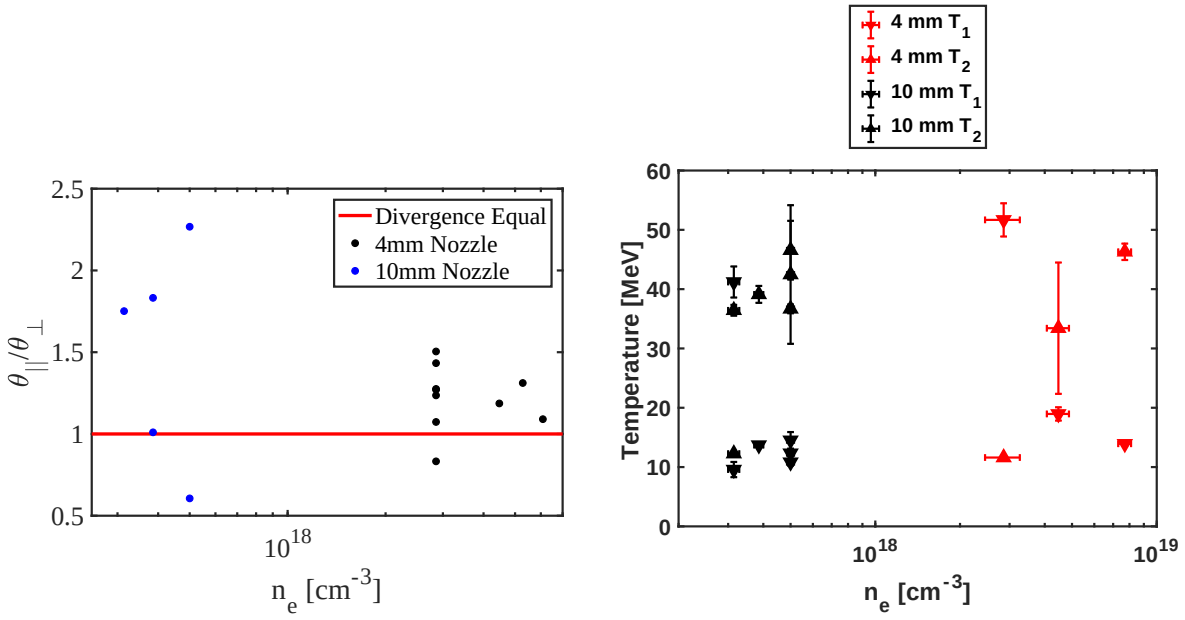

(a) (b)

Figure 3.2: (a) The ratio of the measured electron beam divergence parallel ($\theta_{\parallel}$) and perpendicular ($\theta_{\perp}$) to the linear laser polarization. Ratios greater than 1 indicate an elongation along the laser polarization. (b) All the measured hot (cold), T_1 (T_2) temperatures for the 4 (10) mm gas jet nozzles used in the experiment following the same analysis described in the main text body.

The electron spectra shown in Fig. 3.3 dispersed (a) perpendicular and (b) parallel to the laser polarization were each fit using a single-temperature function ($Ae^{-\frac{E}{T}}$, with amplitude A and effective temperature T in MeV) below and above 100 MeV [Fig. 3.1(b)], yielding low and high temperatures of $T_1 = 19 \pm 0.65$ MeV and $T_2 = 46 \pm 2.45$ MeV, respectively. The perpendicularly dispersed electron signal in (a) (after it is converted to give a linear energy dispersion) shows a clear fork-like structure that begins at electron energies of ∼60 MeV. At 75 MeV, the FWHM of the divergence angle of this fork is 43 mrad [white curve in Fig. 3.3(a)]. The mean total charge contained in this portion of the spectrum (> 60 MeV) is 1.14 ± 0.69 nC. When the electrons are dispersed in the same plane as the laser electric field [Fig. 3.3(b)], no forking structure is seen, and the FWHM beam divergence is instead 21 mrad [red curve in Fig. 3.3(b)] at the same energy. The elliptical beam profile of the electrons shown in Fig. 3.1(c) gives the overall full-angle divergence at half-maximum charge of the electron beam in the two planes as 47 and 27 mrad in the x and y directions, respectively, consistent with the dispersed spectra.

It has been argued that the forking structure gives clear evidence that electrons above 60 MeV are gaining some or most of their energy by the DLA process [136, 1]. Electrons accelerated mainly through DLA generally exhibit higher energy and greater divergence along the laser polarization direction compared to electrons accelerated predominantly through plasma wave acceleration in SM-LWFA. This larger divergence is evident in the forking structure seen only for high-energy electrons dispersed perpendicular to the laser polarization, as in Fig. 3.3(a).

To discern the relative contribution of the various mechanisms to the final energy of the electrons, we simulated the full acceleration process with particle tracking using the quasi-3D algorithm of the OSIRIS PIC simulation framework [28, 30] for laser and plasma parameters similar to those used in the experiment. (Full 3D simulations were prohibitively large for this problem.) As discussed above, this algorithm allows us to unambiguously determine the work done by the longitudinal field of the plasma wave ($E_{z,m=0}$), as well as the transverse ($E_{x,m=1}$)

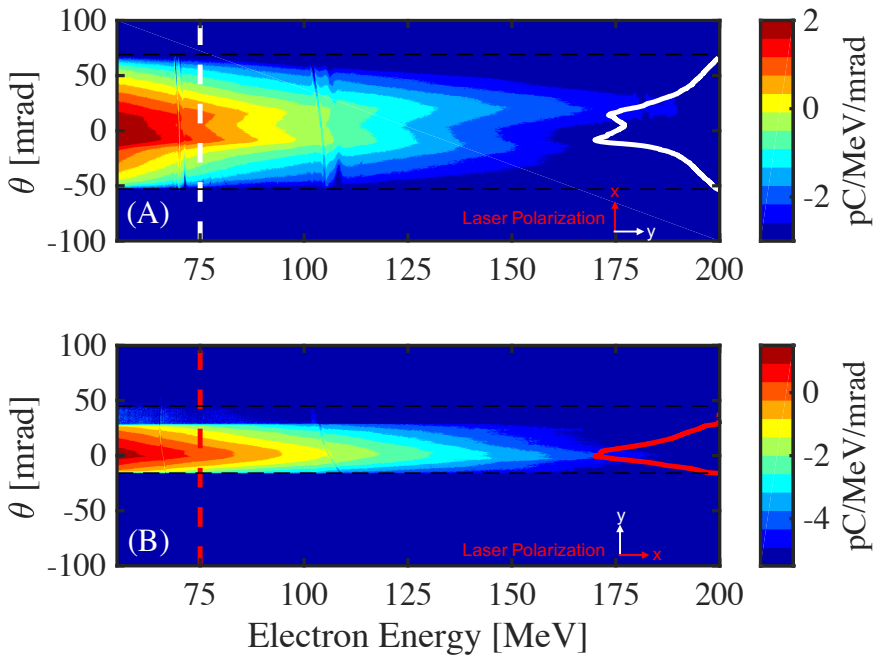

Figure 3.3: Measured electron energy spectra for a plasma with electron density $n_e = 3 \times 10^{17}$ cm^{-3} dispersed (a) perpendicular and (b) parallel to the linear laser polarization direction. The contrast is adjusted and a line-out along the dashed line is plotted (solid line) to emphasize the forking feature in the dispersed electron profile. The dashed black line indicates the acceptance aperture of the magnet. Note that (a) and (b) were taken on two different shots with similar laser energies.

and longitudinal ($E_{z,m=1}$) fields of the laser. This permits a more correct determination of the overall DLA contribution. Here $m = 0$ and $m = 1$ refer to the cylindrical modes corresponding predominantly to the wake and the laser, respectively.

The quasi-3D algorithm uses fields and currents defined on an r-z grid and expanded in azimuthal modes; to simulate LWFA we used modes 0 and 1, where mode 0 (1) mainly captures the wake (laser) fields. In addition, as described later in this chapter, we use a customized field solver that corrects both for dispersion errors of light in vacuum and the time-staggering error of the magnetic field in the Lorentz force [37]. The simulations are carried out in the speed-of-light frame (moving window) with a box of size $95 \times 23.6\, c/\omega_p = 854 \times 212\, \mu$m (the second dimension corresponding only to a half-slice, r starting at 0), where $c/\omega_p = 8.991\, \mu$m for a density of $3.5 \times 10^{17}\, \mathrm{cm}^{-3}$. The number of grid points used was $48000 \times 256 = 1.2 \times 10^7$, with a time step of $\Delta t = 5 \times 10^{-4} \omega_p^{-1} = 30$ fs. The laser pulse has an amplitude of $a_0 = 2.7$, intensity FWHM of 1 ps, spot size of $w_0 = 25.5\, \mu$m and Rayleigh length of $z_R = 1.94$ mm. We use a preformed plasma with a density upramp of $500\, \mu$m followed by a constant-density region, with the laser focused halfway through the upramp. Mobile ions are included along with electrons, with each species having 4 particles per r-z cell and 8 particles in the θ direction, making for a total of 7.9×10^8 particles.

Although the experimental gas jet was 10 mm in length, we find electron energies comparable to those from the experiment after a propagation distance of only 4.68 mm in the simulation. This discrepancy is likely caused by the non-ideal laser spot used in the experiment, while the simulation uses an ideal diffraction-limited Gaussian spatial profile at the laser focus. The non-ideal laser spot could necessitate additional propagation distance for self-modulation and particle trapping to occur in the experiment.

Figure 3.4 shows the envelope of the transverse laser field $E_{x,m=1}$ (green), the plasma density (blue) and the on-axis longitudinal electric field of the plasma wave $E_{z,m=0}$ (red) 4.64 mm into the plasma. Clear modulation of both the laser envelope and plasma waves is evident. However, a hydrodynamic channel (ion motion) is not fully formed (not shown)

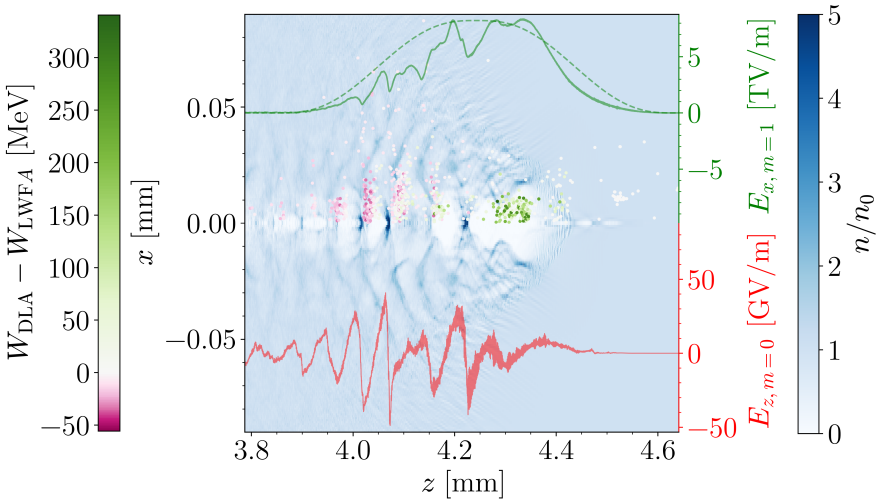

Figure 3.4: Snapshot of the electron density profile after 4.64 mm of propagation (left to right) through the plasma; z and x are the longitudinal and transverse directions, respectively. Also shown are the $m = 0$ longitudinal electric field (SM-LWFA) overlaid in red and the $m = 1$ transverse electric field envelope (DLA) in green. The dashed green line shows the vacuum laser field envelope at the focus. The tracked electrons, with x positions given by their radial distance (only half-space is shown), indicate where in space each acceleration mechanism is dominant. The charge density has been integrated in θ.

within the laser pulse; the ion density remains above $0.9n_0$ across the first bucket (potential well) and above $0.7n_0$ where the laser field is of significant amplitude, where n_0 is the initial plasma density. The wavelength of the plasma wave is increased for the first three buckets by strong beam loading, but for subsequent buckets it is close to $2\pi c/\omega_p$. The plasma electrons trapped by the plasma wave are color-coded to indicate which acceleration mechanism is at work (see subsequent paragraphs). The accelerated electrons group together in the later plasma buckets, where they gain energy predominantly by interacting with the longitudinal field of the wave associated with SM-LWFA. However, the electrons trapped in the front three buckets of the wake gain net energy predominantly through the DLA process as they interact with the peak-intensity portion of the laser pulse (although the SM-LWFA mechanism is also active for some electrons in this region as well). Relativistic self-focusing helps to maintain the peak intensity of the laser pulse (see dashed green line).

To quantify the contribution of each acceleration mechanism (i.e., SM-LWFA and DLA), we use electron tracking in OSIRIS to calculate the work done on each electron by the different spatial components of mode 0 (wake) and mode 1 (laser). Separating the fields by mode clearly shows which longitudinal field component is from the plasma wave ($E_{z,m=0}$) and which is predominantly from the evolving laser field ($E_{z,m=1}$). Without separating the fields by mode, effects from the longitudinal laser electric field can be misattributed to wakefield energy gain or loss—for instance, the total work done on some electrons by the $E_{z,m=1}$ field was $-100\,\mathrm{MeV}$. This energy loss occurs because the longitudinal component of the laser electric field is roughly $\pi/2$ out of phase with the transverse laser electric field, $E_{x,m=1}$. In fact, as discussed previously, the ratio f of average energy gained during betatron oscillation from the parallel ($\langle W_\parallel \rangle$) and perpendicular ($\langle W_\perp \rangle$) fields of a Gaussian laser is given [68] as

$$f \equiv \frac{\langle W_\parallel \rangle}{\langle W_\perp \rangle} = -\frac{2c^2\sqrt{2\gamma}}{\omega\omega_p w_0^2}, \tag{3.5}$$

where w_0 is the beam waist and γ is the gamma factor of a particle. The ratio is negative, indicating that an electron in phase with the transverse laser electric field ($E_{x,m=1}$) loses energy due to the longitudinal laser electric field ($E_{z,m=1}$). Prior to this work, the LWFA and DLA

processes were differentiated by longitudinal and transverse field components, respectively, rather than separated by mode. Consequently the energy loss from the longitudinal laser electric field was often attributed to SM-LWFA electric fields. Experiments and 2D PIC simulations have shown some contribution of the longitudinal laser field to the acceleration of electrons, but for a near-critical-density plasma using foam targets [150], and similar simulation work has shown evidence of energy loss due to the longitudinal laser fields [134].

The work done on each electron is then calculated as follows: $W_{\mathrm{LWFA}} = \int \vec{E}_{m=0} \cdot \vec{v}\, dt$ and $W_{\mathrm{DLA}} = \int \vec{E}_{m=1} \cdot \vec{v}\, dt$. We subtract the work done by each mechanism to obtain a relative energy contribution for each electron (see color of tracked electrons in Fig. 3.4), where a positive (negative) value indicates that the net final energy of the electron is mainly coming from DLA (LWFA). DLA is the dominant energy transfer mechanism for electrons trapped in the front two buckets (Fig. 3.4), whereas SM-LWFA dominates in the later buckets of the plasma wave. As mentioned earlier, due to the low-density plasma (relatively low P/P_c, where P_c is the critical power for relativistic self-focusing [151]), a substantial ion channel—where prior results show DLA dominating the acceleration scheme [72]—does not form within the laser fields.

Figure 3.5 clearly shows that DLA dominates the energy gain for higher-energy electrons where the forking is observed; nearly all of the electrons that make up the forked structure are from the first plasma bucket. The accelerated electrons from OSIRIS are dispersed (a) perpendicular and (b) parallel to the linear laser polarization using a geometry identical to that shown for the experimental results in Fig. 3.3. In this direct comparison to the experimental data, a forking structure is evident only when the electrons are dispersed perpendicular to the laser polarization. Figure 3.5(c) shows that the longitudinal field, $E_{z,m=0}$, of the self-modulated wake (purple dots and curve) dominates energy gain for electrons with energies up to about 40 MeV, at which point the net DLA contribution (red dots and curve) becomes comparable. At around 60 MeV, the dominant acceleration mechanism in the simulation shifts to DLA, with the fork becoming visible at about the

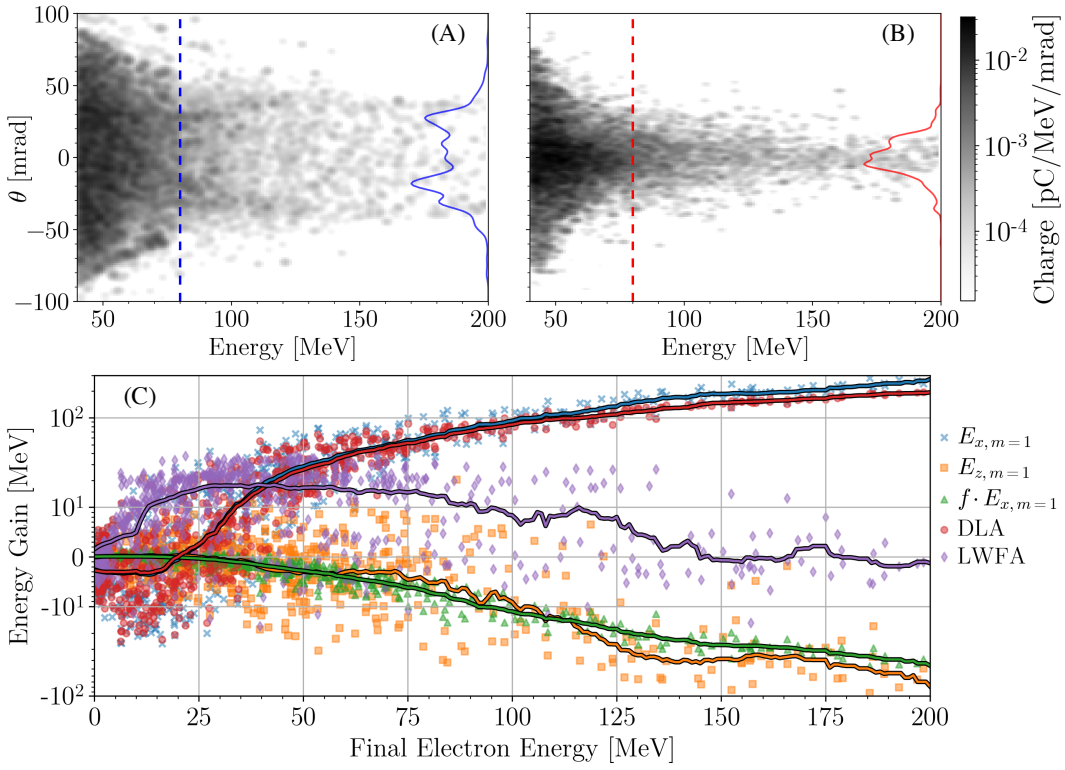

Figure 3.5: Simulated electron spectra dispersed (a) perpendicular and (b) parallel to the linear laser polarization direction. A line-out along the dashed line is plotted (solid line) to emphasize the "horns" of the dispersed beam. (c) The final energy gain due to different field components and mechanisms is shown for numerous tracked electrons (solid lines showing the mean within 20-MeV windows). All data is shown after 4.68 mm of propagation.

same energy for the perpendicularly dispersed electrons both in the experiment [Fig. 3.3(a)] and simulation [Fig. 3.5(a)]. The energy loss from the longitudinal laser field [orange dots and curve in Fig. 3.5(c)] is reasonably approximated for many electrons by the calculation $f \cdot E_{x,m=1}$ shown in green (assumed focused spot size of $w_0 = 19\,\mu$m), which could be used to estimate the energy loss from the longitudinal laser field in simulations where it is difficult to differentiate between longitudinal wake and laser fields (e.g., in 3D or quasi-3D with many azimuthal modes). Regardless, this energy loss is significant and should not be ignored when considering the energy contribution due to DLA.

The simulated electron spectrum is shown alongside the experimental spectra in Fig. 3.1(b), with T_1 and T_2 fits of 14.2 and 45.5 MeV, respectively. Both the number of electrons and the fitted temperatures are similar to the experimental values. The temperature transition, however, occurs near 100 MeV and not 60 MeV (where the DLA mechanism becomes dominant) since the population of electrons accelerated by the DLA mechanism is much smaller than that accelerated by the LWFA mechanism.

3.1.2 Direct laser acceleration for a short-pulse laser in the bubble regime

In the previous section, the combined acceleration from a laser (DLA) and wakefield in a "long" pulse SM-LWFA was examined. In this section, we consider DLA inside the first bucket (bubble) of the short-pulse LWFA regime. This is an example of a regime that cannot be examined using the QSA combined with the PGC as self-trapping, a large amount of photon deceleration, and the breakdown of the PGC occurs. We performed a quasi-3D OSIRIS simulation of the first simulation described in Table 1 of Ref. [152]. The laser has amplitude $a_0 = 53$, spot size $10\,\mu$m, and duration 33 fs and propagates through a 2.5-mm plasma with density $1.5 \times 10^{19}\,\mathrm{cm}^{-3}$. We use 5408×568 cells in z and r and 4 particles in each r-z cell, with 8 particles distributed in ϕ for both electrons and ions. The time step is $0.0075\,\omega_p^{-1}$. This is similar to the simulation in the original bubble regime paper [65]. For these parameters, the pulse length roughly matches half of a plasma period; electrons are continuously injected

and a quasi-monoenergetic beam is formed as the beam dephases in the wake and rotates in phase space [65]. The maximum electron energy reported in Ref. [152] is 3.4 GeV using full-3D PIC simulations.

In order to accurately model phase space rotation, it is important that the the phase velocity of the wake and laser, along with the group velocity of the laser, be modeled accurately. The normal second-order-accurate finite-difference solvers can have significant numerical dispersion errors. Recently we have developed a customized dispersion-free finite-difference solver [37] that corrects for errors in the phase velocity of light propagating through a plasma (affecting both the phase and group velocity through the dispersion relation). The details of this solver are described in Sec. 3.2. With this correction, the group velocity of the laser is closer to the speed of light, and the accelerating length is longer. As a result, we observe a larger maximum energy of 4.7 GeV in a quasi-3D simulation, compared to 4.4 GeV when using the Yee solver with quasi-3D.

In Fig. 3.6 we present plots from various diagnostics for this simulation when the laser has reached 0.07 mm before the 0.1-mm density downramp. We see in (a) the electron density as a function of space (integrated in the azimuthal direction), where we observe total blowout in much of the bubble. A large number of electrons are self-injected near the central axis. In (b) we show the longitudinal $m = 0$ electric field corresponding to the wakefield in blue and green. Layered on top of this field is the $m = 1$ transverse field corresponding to the laser. We see the main laser field is located at the front of the bubble, but that fields of significant amplitude are also present at the back of the bubble. These fields drift backward from the main pulse due to photon deceleration, then are kept at the back of the bubble (and move forward) due to photon acceleration [49], where they can accelerate electrons through the DLA mechanism.

We observe in Fig. 3.6(c) the main body of accelerated particles between 1780 and 1820 c/ω_p, which are rotating in phasespace due to their progression through the wake. This occurs when particles at the front of the beam enter the decelerating phase of the wake.

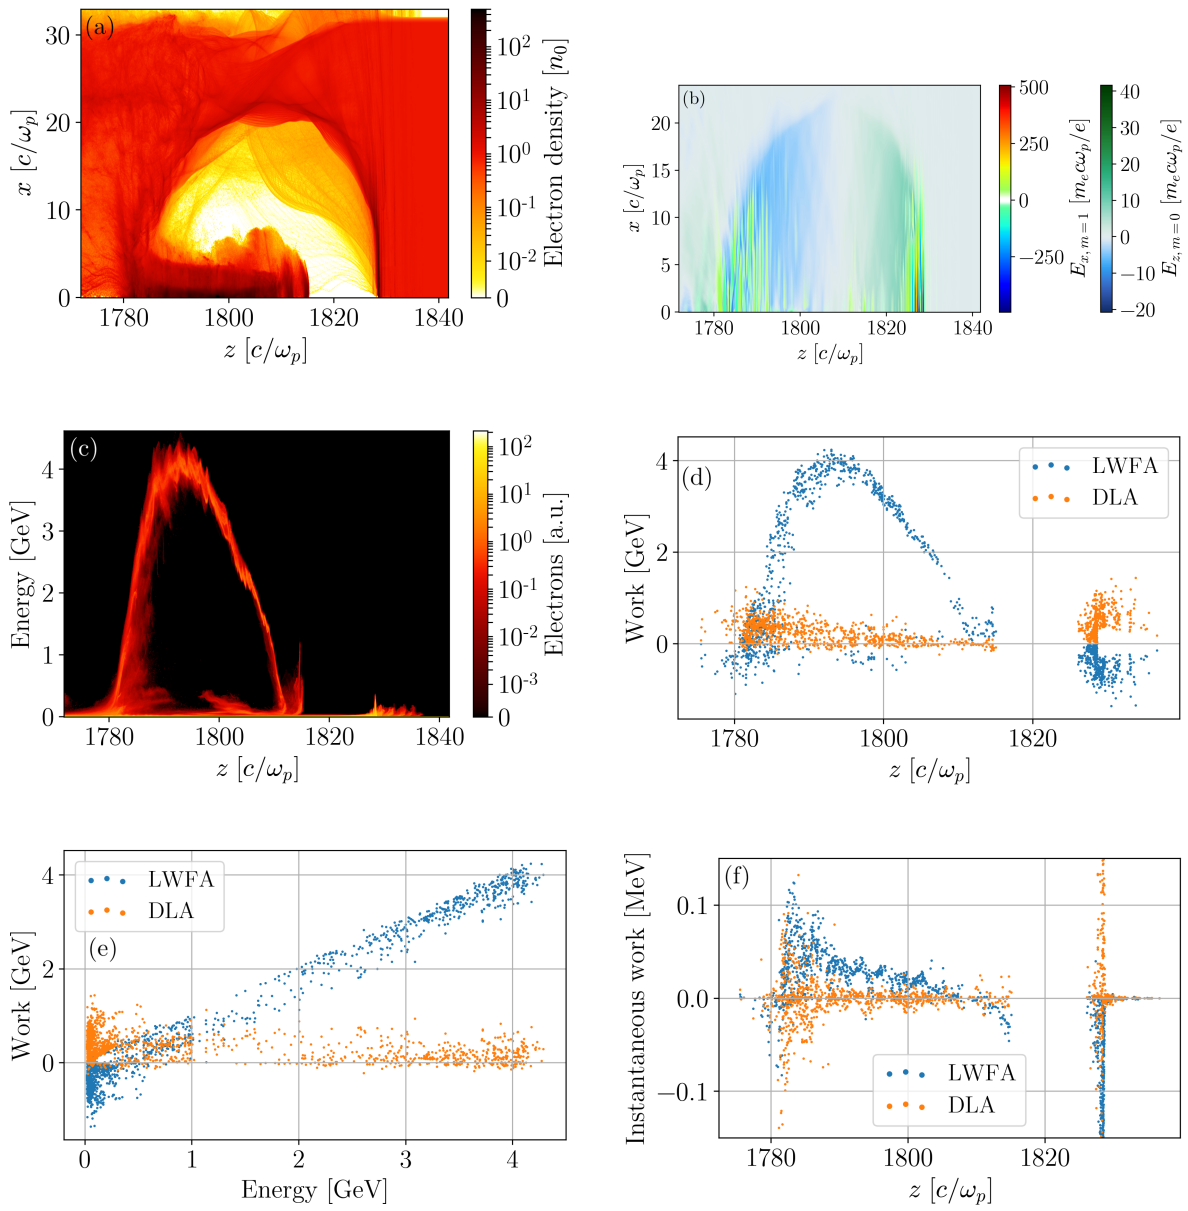

Figure 3.6: OSIRIS simulation data from run with $a_0 = 53$ as described in the body text, just before the density downramp. (a) Electron density. (b) Laser and wake fields in the bubble. (c) Energy–space phasespace showing the rotation of particles. Total work done by LWFA and DLA processes (d) as a function of space and (e) as a function of energy. (f) Instantaneous work done over two time steps as a function of space.

However, we also note a group of particles at the front between 1820 and 1840 c/ω_p, located near the peak of the laser pulse. In (d) we observe the cumulative work done by the LWFA and DLA processes on tracked particles as a function of space. The particles at the front of the bubble receive all their energy from DLA. We also observe that particles between 1780 and 1810 c/ω_p exhibit positive DLA, with this contribution diminishing towards the front of the bubble. This is because the on-going process of photon deceleration gradually increases the laser energy located at the back of the bubble. Thus the electrons that were injected first (now located near the front-middle portion of the bubble) experienced only wakefield acceleration, and electrons injected later in time (located at smaller values of z) experience increasing contributions from DLA.

In Fig. 3.6(e) we observe the cumulative work done by the LWFA and DLA processes as a function of particle energy. We see that particles of many different energies receive significant energy from DLA, and that rarely is the energy contribution negative. Finally, in (f) we see the instantaneous work (over two time steps) done on electrons as a function of space. We can see that only at the back of the bubble (1780–1790 c/ω_p) is the DLA mechanism active, corresponding to the same region where we see laser fields in (b). In addition, the strength of the wakefield drops off linearly moving towards the front of the bubble.

These results show that energetic electrons can receive a significant contribution of their energy from DLA, even for an LWFA operating in what is referred to as the bubble regime. When comparing the DLA energy calculations using the $m = 0$ and $m = 1$ modes as in Eq. (3.4) to those using only the longitudinal and transverse components as in Eq. (3.3), it is clear that the traditional component method overestimates the DLA energy contribution by an average of 50% for particles with energies over 2 GeV. This error occurs because of neglecting the energy loss from the longitudinal component of the laser, which is properly taken into account when grouping modes.

To verify that the quasi-3D simulation is accurately capturing the physics, we also perform a full three-dimensional simulation with 3775×278 cells (note that the transverse simulation

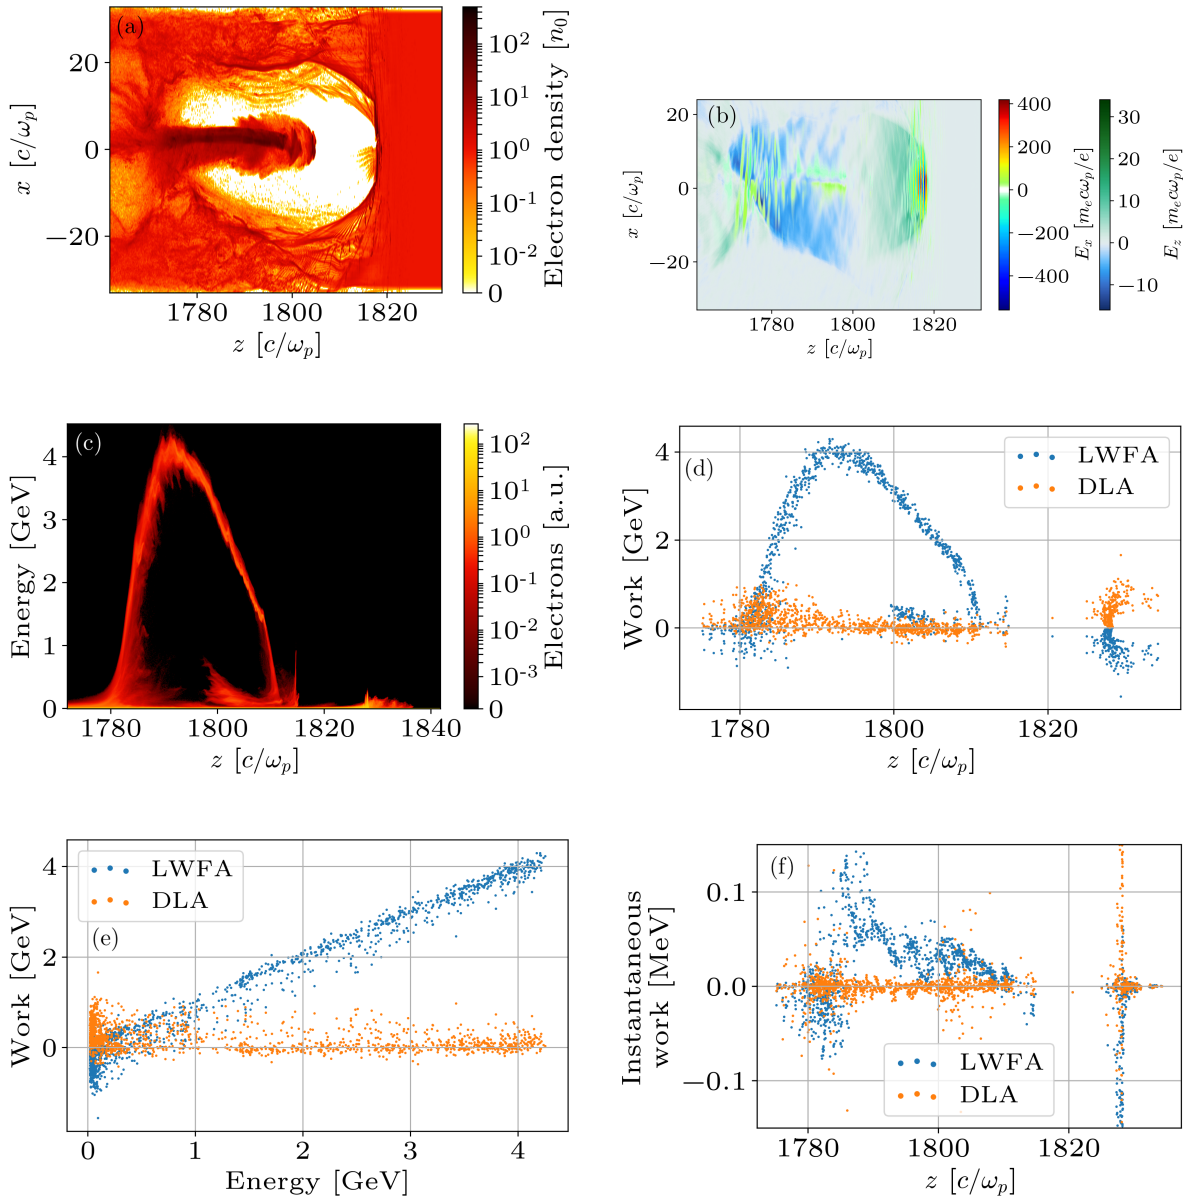

Figure 3.7: Three-dimensional OSIRIS simulation data from run with $a_0 = 53$ as described in the body text, just before the density downramp. (a) Electron density. (b) Laser and wake fields in the bubble. (c) Energy–space phasespace showing the rotation of particles. Total work done by LWFA and DLA processes (d) as a function of space and (e) as a function of energy. (f) Instantaneous work done over two time steps as a function of space.

size is twice as big in 3D) and time step $0.012516\,\omega_p^{-1}$, but otherwise identical numerical and physical parameters. The results are shown in Fig. 3.7, in the same manner as for Fig. 3.6. Most trends are identical to the quasi-3D simulation, but we observe a smaller DLA energy gain for high-energy electrons in three dimensions (note that in Cartesian geometry we cannot compute the DLA contribution as accurately with mode separation).

3.2 Customized field solver for more accurate simulation of direct laser acceleration

PIC simulations have long been used to study DLA of electrons in underdense plasmas [70, 78, 136, 153, 134, 74]. Most PIC codes use the Boris push [25] and Yee finite-difference field solver [24], which both have very advantageous properties, have been well analyzed, and can provide high-performance software. For example, the Boris algorithm conserves energy exactly for particle motion in a magnetic field and also conserves phase space volume [23]. The Yee mesh staggers electric and magnetic fields in time, yielding increased accuracy for a given time step. However, here we wish to discuss two errors from the Yee field solver that negatively impact PIC simulations, particularly those of DLA. Many of the arguments and results presented in this section are described in Ref. [37].

3.2.1 Errors in numerical dispersion

The Yee mesh discretizes space and time, yielding essentially a unique medium through which light waves propagate. Here we examine the dispersion relation of light in this medium to determine the properties of its propagation. First we give the discrete forms of Faraday's and Ampere's laws as

$$\begin{aligned} d_t \boldsymbol{B} &= -\boldsymbol{d}_E \times \boldsymbol{E}, \\ d_t \boldsymbol{E} &= \boldsymbol{d}_B \times \boldsymbol{B} - \boldsymbol{J}, \end{aligned} \tag{3.6}$$

where d_t and $\boldsymbol{d}_{E,B}$ are generalized finite-difference operators. In this section we use normalized units in which time and length are normalized to reciprocals of arbitrary frequency ω_n and wavenumber $k_n = \omega_n/c$, charge to e, mass to m_e, density to $n_n = \omega_n^2 \varepsilon_0 m_e/e^2$, and current density to $en_n c$. The electric and magnetic fields are then normalized to $m_e c\omega_n/e$. Note that the spatial finite-difference operators acting on $\boldsymbol{E}$ and $\boldsymbol{B}$ can be different, but that $\boldsymbol{d}_B$ should be assumed in the continuity equation for a charge-conserving scheme where Gauss's law is maintained (or used to directly solve Gauss's law). This means that for Gauss's law we have $\boldsymbol{d}_B \cdot \boldsymbol{E} = \rho$.

Performing a Fourier transform in time and space of Eq. (3.6) yields

$$\begin{aligned} [\omega]_t \tilde{\boldsymbol{B}} &= [\boldsymbol{k}]_E \times \tilde{\boldsymbol{E}}, \\ [\omega]_t \tilde{\boldsymbol{E}} &= -[\boldsymbol{k}]_B \times \tilde{\boldsymbol{B}} - i\tilde{\boldsymbol{J}}, \end{aligned} \tag{3.7}$$

where an overtilde represents a Fourier transform, and $[\omega]_t$ and $[\boldsymbol{k}]_{E,B}$ are the counterparts of the discrete finite-difference operators in Fourier space. The bracket $[\omega]$ and $[\boldsymbol{k}]$ operators are functions of ω and $\boldsymbol{k}$ as well as the time step and cell size. They approach the continuous limit as the time step and cell size approach zero. Here the electric and magnetic fields are assumed to have the behavior $e^{i(\omega t - \boldsymbol{k}\cdot\boldsymbol{r})}$, where ω and $\boldsymbol{k}$ are the frequency and wavenumber of an electromagnetic wave. We now choose to view the light wave to be propagating in vacuum, meaning that we can neglect the source terms $\boldsymbol{J}$ and ρ; this gives Gauss's law as $i[\boldsymbol{k}]_B \cdot \tilde{\boldsymbol{E}} = 0$. The numerical dispersion relation can then be obtained in vacuum as

$$[\omega]_t^2 - [\boldsymbol{k}]_B \cdot [\boldsymbol{k}]_E = 0. \tag{3.8}$$

Further assuming the light wave to be propagating only in the $\hat{1}$-direction yields a simplified numerical dispersion relation,

$$[\omega]_t^2 - [k]_{B1}[k]_{E1} = 0. \tag{3.9}$$

For the standard Yee mesh where the electric field is defined a half time step removed from the magnetic field, the operator $[\omega]_t$ has the form $[\omega]_t = \sin\left(\frac{\omega\Delta t}{2}\right)/\frac{\Delta t}{2}$, where Δt is the

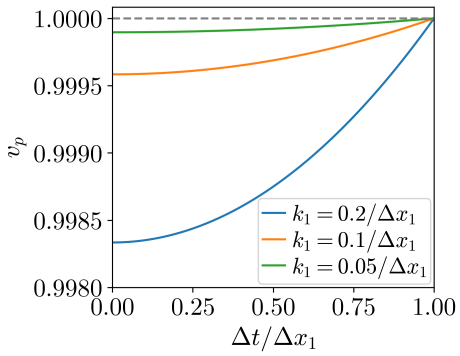

Figure 3.8: Phase velocity of light waves with three different wavenumbers as a function of simulation time step for the standard second-order Yee solver. Note that all waves travel slower than the speed of light for any time step less than the one-dimension CFL limit, $\Delta t \leq \Delta x_1$.

time step chosen for a particular simulation. Plugging this into Eq. (3.9) yields the phase velocity as

$$v_p = \frac{\omega}{k_1} = \frac{2}{k_1 \Delta t} \arcsin\left(\frac{\Delta t}{2}\sqrt{[k]_{B1}[k]_{E1}}\right). \tag{3.10}$$

The electric and magnetic fields are also staggered in space on the standard Yee mesh, and the finite-difference operators of arbitrary order, p, acting on $\boldsymbol{E}$ and $\boldsymbol{B}$ have the form

$$[k]_{B1,E1} = \sum_{j=1}^{p/2} C_j^{B,E} \frac{\sin\left[(2j-1)\frac{k_1 \Delta x_1}{2}\right]}{\Delta x_1/2}, \tag{3.11}$$

where $C_j^{B,E}$ is the stencil coefficient and Δx_1 is the chosen simulation resolution size in the $\hat{1}$-direction. The standard second-order Yee solver uses the same solver stencil for both Faraday's and Ampere's equations and has $[k]_{B1} = [k]_{E1} = \sin\left(\frac{k_1 \Delta x_1}{2}\right) / \frac{\Delta x_1}{2}$. We can thus rewrite the dispersion relation for the standard second-order Yee solver as

$$v_{p,\text{Yee}} = \frac{2}{k_1 \Delta t} \arcsin\left[\frac{\Delta t}{\Delta x_1} \sin\left(\frac{k_1 \Delta x_1}{2}\right)\right]. \tag{3.12}$$

The mixing of Δt and Δx_1 in Eq. (3.12) is an important feature of the Yee solver. For example, when doing a convergence test, we often wish to decrease the time step to ensure

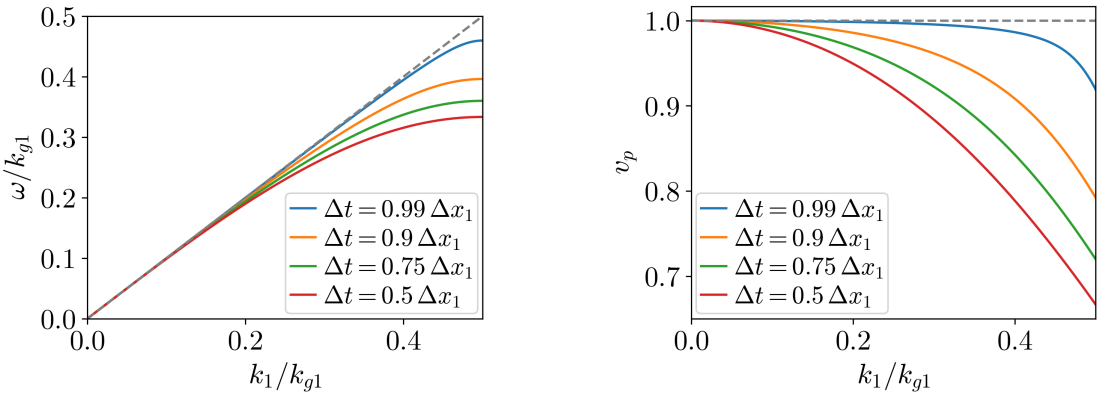

(a) (b)

Figure 3.9: (a) Frequency and (b) phase velocity for various values of simulation time step as a function of wavenumber using the standard second-order Yee solver. Here $k_{g1} = 2\pi/\Delta x_1$.

accuracy of the simulation. The particle push is always more accurate when the time step is decreased. However, in Fig. 3.8 we show the phase velocity for the standard Yee solver as given by Eq. (3.12) for light waves with various wavenumbers as a function of the time step. A wave with $k_1 = 0.2/\Delta x_1$ has $2\pi/0.2 \approx 31$ points per wavelength, which is reasonable for many applications. We see from the figure that the phase velocity of the light wave deviates from c for any time step below the Courant–Friedrichs–Lewy (CFL) stability condition for a one-dimensional simulation, $\Delta t \leq \Delta x_1$. Thus in order to maintain proper (or the same) dispersion during a convergence test, both the spatial and temporal discretization steps must be decreased in the process, resulting in very computationally expensive runs in higher dimensions.

Additionally, in Fig. 3.9 we show the (a) frequency and (b) phase velocity for various values of simulation time step as a function of wavenumber, where $k_{g1} = 2\pi/\Delta x_1$. We note that in comparison of Fig. 3.8, one light wave of interest has $k_1 = 0.2/\Delta x_1 = 0.2k_{g1}/2\pi \approx 0.03k_{g1}$. We see that most high-$k$modes travel at significantly less than the speed of light regardless of time step, which is one reason that current filters are widely used in electromagnetic PIC

codes. For any PIC simulation, spatial and temporal discretization steps must be chosen such that the physics of interest is not only well-resolved, but exhibits proper dispersion. To see dispersion properties of Yee-type field solvers of higher orders, see Fig. 1 of Ref. [37].

3.2.2 Inaccurate calculation of the Lorentz force

We mentioned previously that the Boris algorithm conserves energy for a particle moving in the presence of a magnetic field (electric field absent). However, the staggering of $\boldsymbol{E}$ and $\boldsymbol{B}$ in time on the Yee mesh itself introduces a spurious force exerted on particles moving in a laser field when using the standard Yee field solver. In this section, we assume the laser to be polarized in the $\hat{2}$-direction (still traveling in the $\hat{1}$-direction). A particle with charge q is pushed according to

$$\frac{\gamma^{n+1/2}\beta_2^{n+1/2} - \gamma^{n-1/2}\beta_2^{n-1/2}}{\Delta t} = q\left(E_2^n - \bar{\beta}_1^n \times \bar{B}_3^n\right), \tag{3.13}$$

where the laser fields E_2 and B_3 are interpolated from the spatial grid points, and β_i refers to the velocity of the particle in the $\hat{i}$-direction. The superscript x^n represents the quantities at $t = n\Delta t$, and the overbar $\bar{x}$ represents interpolation in time. In the standard PIC algorithm, the time interpolation is usually fulfilled by a simple average, i.e., $\bar{x}^n = \left(x^{n-1/2} + x^{n+1/2}\right)/2$. Since the magnetic field components are defined on the half time step, whereas the Lorentz force is evaluated at the integer time step, B_3 needs to be interpolated in time.

Consider a single particle to be moving at very near the speed of light in the $\hat{1}$-direction within these laser fields. Ideally—with perfect dispersion and without the need to perform time-averaging of the magnetic field—we would have according to Faraday's law that $B_3^n = E_2^n$. The total force on the particle would then be $qE_2^n\left(1 - \bar{\beta}_1^n\right) \approx 0$. However, because of the time-averaging that is necessary in Eq. (3.13), $\bar{B}_3^n$ does not equal E_2^n with sufficient precision for such particles moving near the speed of light in the $\hat{1}$-direction. We seek to quantify and eliminate this error by our choice of field solver. It should be noted that although the field components are also spatially staggered on the Yee mesh, our derivation shows that the

spatial staggering has no contribution to the spurious Lorentz force.

From the discretized Maxwell equations in Eq. (3.7), we have that

$$\tilde{B}_3 = \frac{[k]_{E1}}{[\omega]_t}\tilde{E}_2, \tag{3.14}$$

from which it can be shown [37] that in two dimensions the transverse Lorentz force exerted on the particle is

$$\frac{\tilde{F}_2(\omega, k_1, k_2)}{q} = \left[\tilde{E}_2(\omega, k_1, k_2) - \bar{\beta}_1 \tilde{B}_3(\omega, k_1, k_2)\cos\frac{\omega\Delta t}{2}\right] \times \tilde{S}(-k_1, -k_2), \tag{3.15}$$

where $\tilde{S}$ is the Fourier transform of the interpolation function. The factor of $\cos\frac{\omega\Delta t}{2}$ is due to the time staggering and corresponding average required in Eq. (3.13), but we can see that the spatial staggering has no impact. Using Eqs. (3.9) and (3.14)in Eq. (3.15) yields

$$\frac{\tilde{F}_2}{q} = \tilde{E}_2\left[1 - \bar{\beta}_1\sqrt{[k]_{E1}/[k]_{B1}}\cos\frac{\omega\Delta t}{2}\right]\tilde{S}. \tag{3.16}$$

For the standard PIC algorithm with $[k]_{E1} = [k]_{B1}$, the factor $\cos\frac{\omega\Delta t}{2}$ cannot be eliminated from Eq. (3.16), resulting in an incorrect force on a particle moving within laser fields in vacuum.

3.2.3 Improved Maxwell solver with dual k-space operators

The combination of the errors of numerical dispersion and a spurious Lorentz force described in Secs. 3.2.1 and 3.2.2, respectively, are especially detrimental to simulations of DLA. In such DLA cases, electrons have long-time interactions with the laser fields in very low-density regions. Errors in the phase velocity of the laser or the force it imparts on the particles can result in large differences in final energy. We seek to find a finite-difference solver that both improves the numerical dispersion relation and compensates for the spurious force induced by the time stagger of the $\boldsymbol{E}$ and $\boldsymbol{B}$ fields. We accomplish this by determining $[k]_{E1}$ and $[k]_{B1}$ operators (in Fourier space) that minimize or eliminate both of these errors, and then to develop finite-difference operators (in real space) that provide those desired $[k]_{E1}$ and

$[k]_{B1}$ operators. If one is using an FFT-based algorithm, the desired operators can be used in Fourier space directly. This is commonly done for pseudo-spectral analytic time-domain solvers, as described for some in Ref. [154].

To minimize the dispersion errors, the $[k]_{E1}$ and $[k]_{B1}$ operators should be properly selected so that the numerical version of the ratio k_1^2/ω^2 approximates the real physics as much as possible:

$$\frac{[k]_{E1}[k]_{B1}}{[\omega]_t^2} \to \frac{k_1^2}{\omega^2}. \tag{3.17}$$

Furthermore, to minimize spurious terms in the Lorentz force, we would like to simultaneously find $[k]_{E1}$ and $[k]_{B1}$ to satisfy

$$1 - \bar{\beta}_1 \sqrt{[k]_{E1}/[k]_{B1}} \cos\frac{\omega\Delta t}{2} \to 1 - \bar{\beta}_1. \tag{3.18}$$

The first condition, Eq. (3.17), is satisfied by setting $[k]_{E1}[k]_{B1} = [k]_{1,t}^2$, where

$$[k]_{1,t} \equiv \sin\left(\frac{k_1\Delta t}{2}\right) / \frac{\Delta t}{2}. \tag{3.19}$$

(Here we remember that $[\omega]_t = \sin\left(\frac{\omega\Delta t}{2}\right) / \frac{\Delta t}{2}$ for the standard PIC algorithm.) Substituting this relationship for the $[k]$ terms into the second condition, Eq. (3.18), then gives

$$\begin{aligned} [k]_{E1} &= \frac{[k]_{1,t}}{\cos\left(\frac{\omega\Delta t}{2}\right)}, \\ [k]_{B1} &= [k]_{1,t} \cos\left(\frac{\omega\Delta t}{2}\right). \end{aligned} \tag{3.20}$$

It is difficult to ensure that these equations are simultaneously satisfied in general due to the ω-dependence inside the cosine terms. However, we are interested in ensuring that this relationship is satisfied for light waves propagating in tenuous plasmas where $\omega \approx k_1$; therefore we replace $\omega = k_1$ in Eq. (3.20) to obtain

$$\begin{aligned} [k]_{E1} &= \frac{[k]_{1,t}}{\cos\left(\frac{k_1\Delta t}{2}\right)} = \frac{2}{\Delta t}\tan\left(\frac{k_1\Delta t}{2}\right), \\ [k]_{B1} &= [k]_{1,t}\cos\left(\frac{k_1\Delta t}{2}\right) = \frac{1}{\Delta t}\sin\left(k_1\Delta t\right). \end{aligned} \tag{3.21}$$

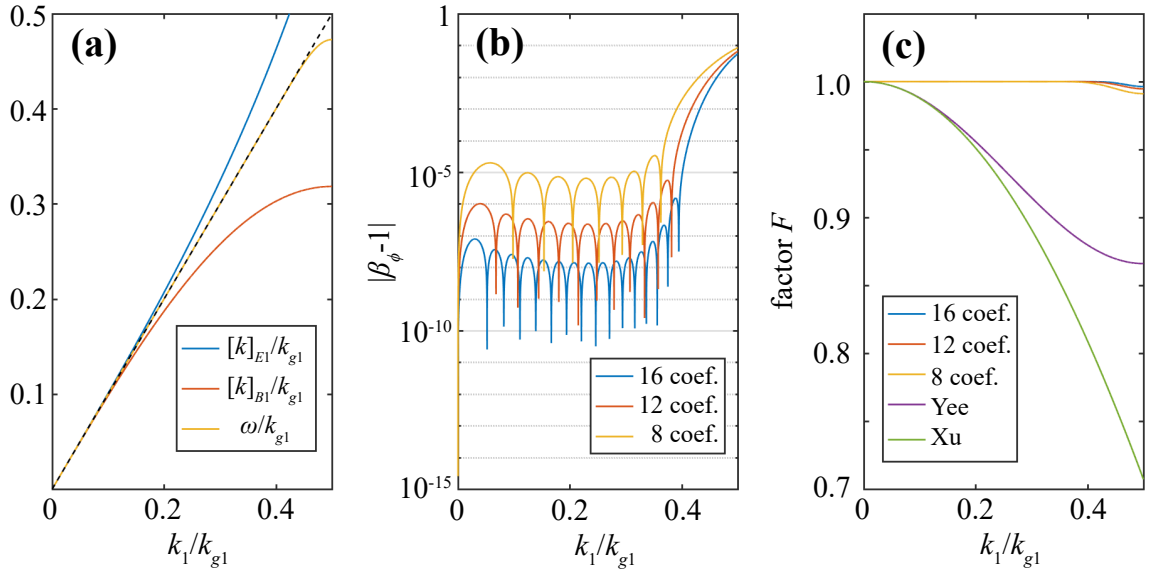

Figure 3.10: (a) Calculated $[k]_{E1}$ and $[k]_{B1}$ operators fitted by the coefficient customization method (using a 16-coefficient stencil). The yellow line represents the numerical dispersion relation, $[\omega]_t = \sqrt{[k]_{E1}[k]_{B1}}$. (b) Error in phase velocity for solver stencils with varying number of coefficients. (c) Comparison of the cancellation factor $F \equiv \sqrt{[k]_{E1}/[k]_{B1}}\cos\frac{\omega\Delta t}{2}$ (ideally $F = 1$) between the proposed solvers with different stencils, the standard Yee solver and the Xu solver. The numerical parameters are $\Delta x_1 = 0.2k_0^{-1}$, $\Delta t = 0.4\Delta x_1$ and $k_{g1} \equiv 2\pi/\Delta x_1$.

The operators in Eq. (3.21) are approximated for a finite-difference method in broad regions of k_1-space via the methodology in Ref. [155], where the solver stencil is extended and customized coefficients are used. In comparison to the phase velocity for the Yee solver in Eq. (3.12) that mixes Δt and Δx_1, we observe only the dependence on Δt in Eq. (3.21). Since the dispersion relation $\omega = k_1$ is satisfied for any value of time step, one can decrease the time step to well below the CFL limit (which makes the pusher more accurate) and not suffer the dispersion errors pictured in Fig. 3.8.

We note that choosing a solver with $[k]_{E1} = [k]_{B1} = [k]_{1,t}$ is exactly the solver proposed in Ref. [36] to reduce field errors surrounding relativistic particles, which we refer to herein as

the "Xu" solver. The scheme proposed here, which we refer to herein as the "Fei" solver, can be viewed as a time-stagger-correction version of the Xu solver. We show in Fig. 3.10(a) the calculated $[k]_{E1}$ and $[k]_{B1}$ operators using a 16-coefficient stencil, along with the resulting numerical dispersion relation given by $[\omega]_t = \sqrt{[k]_{E1}[k]_{B1}}$. The errors in phase velocity, $|v_p - 1|$ (shown as $|\beta_\phi - 1|$ in the figure) are shown in (b) for various numbers of coefficients. Using 16 coefficients, the phase velocity errors are below 10^{-7} for a broad range of wavenumbers. A low-pass filter can be used to remove the high-k modes. Finally, in (c) we show the factor multiplying the particle velocity in the Lorentz force equation, $F \equiv \sqrt{[k]_{E1}/[k]_{B1}}\cos\frac{\omega\Delta t}{2}$, for a variety of solvers and number of coefficients. This factor should be equal to unity in the ideal case, and we see it deviates significantly for all solvers except for the Fei solver. In particular, the error is significant even for well-resolved modes ($k_1 \approx 0.03 k_{g1}$) using the standard Yee solver. For more details on the properties, implementation and performance of the proposed Fei solver, see Ref. [37].

3.2.4 Comparison of field solvers for full-scale, self-consistent particle-in-cell simulations of direct laser acceleration

Using the field solver proposed in Sec. 3.2.3, we now wish to characterize its impact on a full-scale, three-dimensional PIC simulation of an LWFA with DLA. The scenario we will be simulating was originally reported in Ref. [1], which compares experiment and simulation for a laser in the quasi-blowout regime where the laser driver overlaps the trapped electrons in the first bucket. In that work, the experimental electron spectra dispersed in energy perpendicular to the laser polarization showed forking at high energies, a clear sign of the presence of DLA. The experimental data from that work is reproduced here for comparison in Fig. 3.11.

The three-dimensional OSIRIS simulation performed in Ref. [1] had $a_0 = 2.03$, $\lambda = 815\,\text{nm}$, a laser spot size of $w_0 = 6.7\,\mu\text{m}$, plasma density $n_e = 1.43 \times 10^{19}\,\text{cm}^{-3}$ from a 99.9% He and 0.1% N_2 neutral mix, and FWHM of laser intensity of $\tau = 45\,\text{fs}$. The grid was $1814 \times 320 \times 320$

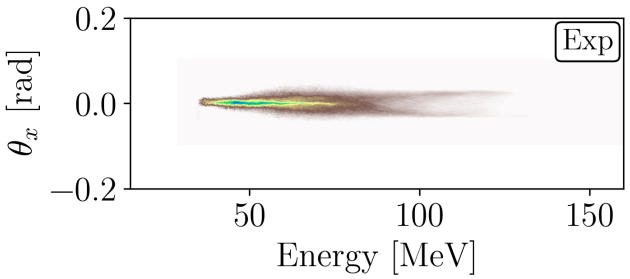

Figure 3.11: Experimental electron spectra from Ref. [1] dispersed perpendicular to the laser polarization direction (reproduced here). The parameters for this show are gas cell length of 900 μm, $n_e = 1.4 \times 10^{19}\,\mathrm{cm}^{-3}$, $a_0 = 1.9$ and $\frac{\omega_p \tau}{2\pi\sqrt{a_0}} = 1.1$.

cells with $2 \times 2 \times 2$ particles per cell, and $k_0 \Delta z = 0.209$ and $k_p \Delta x, \Delta y = 0.120$. The laser was linearly polarized in the x-direction (propagating in the z-direction) and was focused halfway up a 100-μm density upramp, then propagated through a 430-μm constant-density region before exiting through a 150-μm density downramp. The simulations showed that the electron spectra exhibited a forking structure along the direction parallel to the laser polarization, and the tracking of individual electrons revealed that those accelerated to the highest of energies received the majority of their energy from DLA.

After receiving from the author of Ref. [1] the original input file to the OSIRIS simulation described above, we have redone that simulation using the Yee, Xu and Fei field solvers in three dimensions. We have also performed the simulation in the quasi-3D geometry with the Fei field solver, which completes in roughly an order or magnitude less wall clock time with four times fewer processors compared to full 3D for this particular simulation. In Fig. 3.12 we show the electron spectra on an angle versus energy plot, where θ_x is defined with respect to the laser propagation direction in the plane containing the direction of the laser polarization. This simulation in (a) uses identical field solvers and particle pushers as was used in the simulations of Ref. [1]. Comparing (a) with Fig. 4(a) of Ref. [1], one can see there are slight differences between the results. This is because of various changes and improvements to the algorithms in the base OSIRIS software over a number of years. However, the majority of the

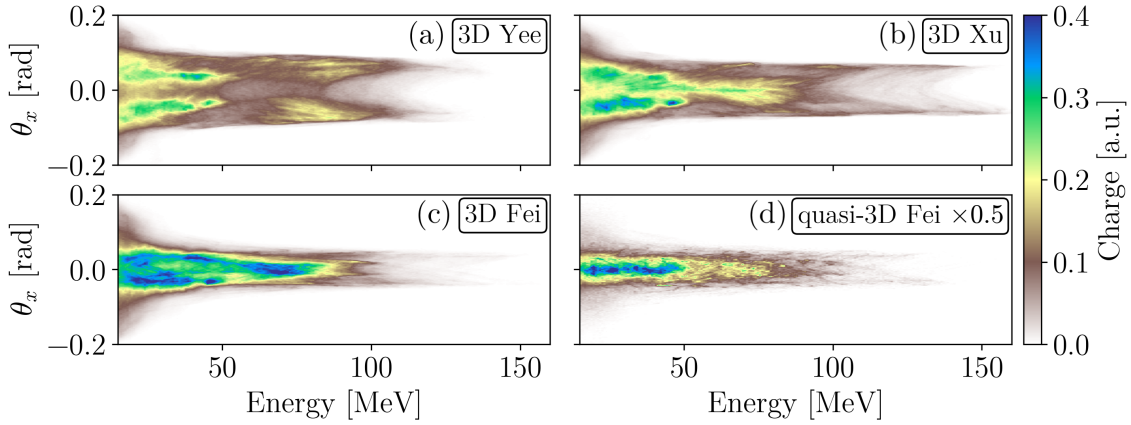

Figure 3.12: Electron spectra from four simulations shown as angle (parallel to laser polarization) as a function of energy, with the data from the quasi-3D simulation multiplied by 0.5 for visibility.

trends remain the same.

The angular spread in Fig. 3.12 is widest for the Yee solver, is narrower for the Xu solver, and is narrowest for the Fei solver. In addition, the Yee solver shows no on-axis group of particles at the highest energies, with almost all energetic electrons exhibiting some finite angle. Comparing the simulation data against the experimental data in Fig. 3.11 reveals that using the Fei solver provides the most accurate representation for the particle energy, narrow divergence of the energy distribution, and a forking at high energy. Note that any particles below 40 MeV were not captured by the experimental spectrometer. We also point out the similarity between the full-3D and quasi-3D simulations when using the Fei solver.

Next, we examine the behavior of the accelerated bunch in real space, shown in Fig. 3.13 for the various solvers. We see large differences between each field solver, with perhaps the most significant difference being the transverse displacement of the beam. For the Yee solver in (a), particles actually drift into the bubble sheath; we also see that the distribution of the beam in the z coordinate is approximately sinusoidal, with particle bunches clustered at the extremes in x. This distribution agrees with the large angular spread seen in Fig. 3.12(a).

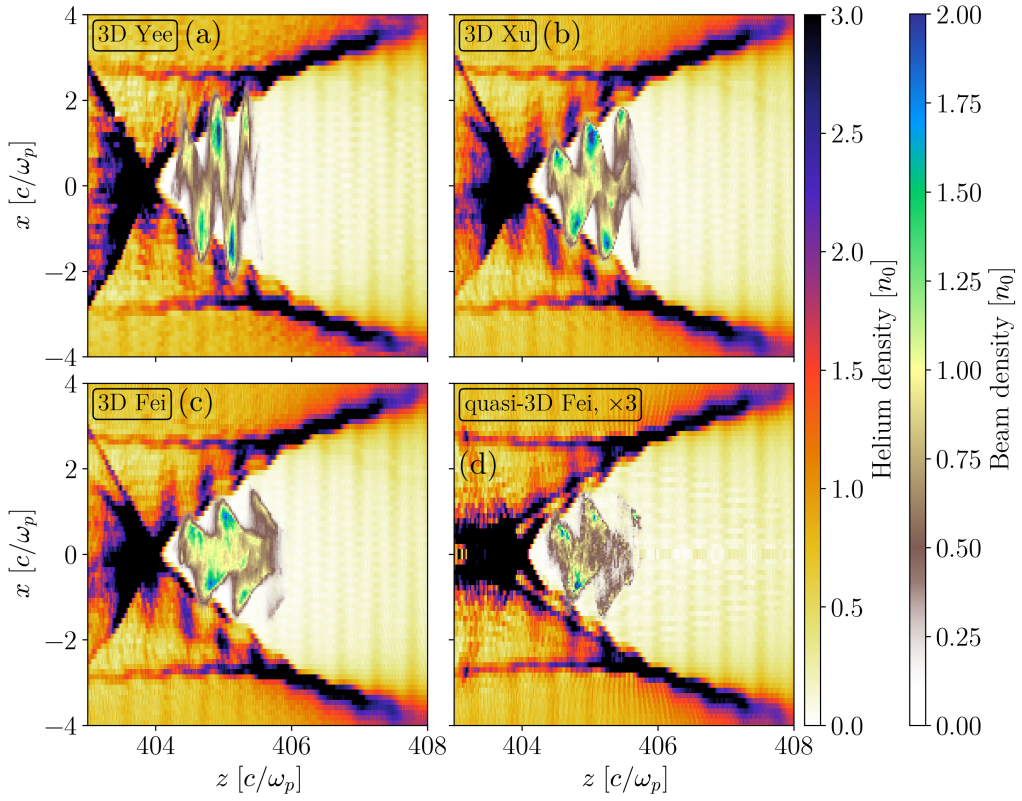

Figure 3.13: Position of accelerated electron beam for four simulations, with the data from the quasi-3D simulation multiplied by 3 for visibility.

For the Xu solver in Fig. 3.13(b), the particles extend right up to the sheath but have not penetrated through it by this time. We see that the transverse motion is less than in (a), that some particles are clustered near the axis, and that the distribution in z is less sinusoidal. Finally, in (c) we see that the particles reside farther from the sheath and have a different overall shape. Although not obvious from the PIC simulation, we show in Sec. 3.3.2—by integrating the equations of motion for electrons executing betatron motion with similar laser and plasma parameters—that the beam shape in (c) is most correct among the three field solvers [see Fig. 3.16(c)]. Finally, we observe that the beam shape in (d) is very similar to the full-3D simulation in (c), but the particle statistics are not as good.

We reran each simulation while tracking 1,500 individual electrons. With the tracking data, we can evaluate the fields acting on each particle at a given time step to determine the work done on the particle. In three dimensions, the standard option for doing this is to use the longitudinal (E_1 or E_z) and transverse (E_2 or E_x) fields for the LWFA and DLA processes, respectively, as in Eq. (3.3). However, we can also attempt to include a correction for the energy loss from the longitudinal component of the laser by adding the quantity $f * E_1$ to W_{DLA} and subtracting it from W_{LWFA}, where f is defined for a Gaussian laser in Eq. (3.2). In Fig. 3.14(a)–(c) we show the calculated work from LWFA and DLA on tracked electrons as a function of energy for the three different solvers in three dimensions. In (d) we apply the longitudinal laser field correction to the three-dimensional simulation using the Fei solver, which greatly lowers the energy gain from DLA for high-energy electrons. Note that the calculation of f assumes a Gaussian laser in free space and requires estimating the spot size of the laser, which does not completely describe the laser fields at the back of the bubble.

In Fig. 3.14(e) and (f) we show the calculated work using the transverse [Eq. (3.3)] and modal [Eq. (3.4)] methods in quasi-3D, respectively, where (f) is the most accurate computation. We see that in comparison to (f), (d) with the applied correction is much more accurate than (c), which overestimates the DLA energy contribution. We also note that the DLA energy gain is lower overall when using the Yee solver due to the improper dispersion

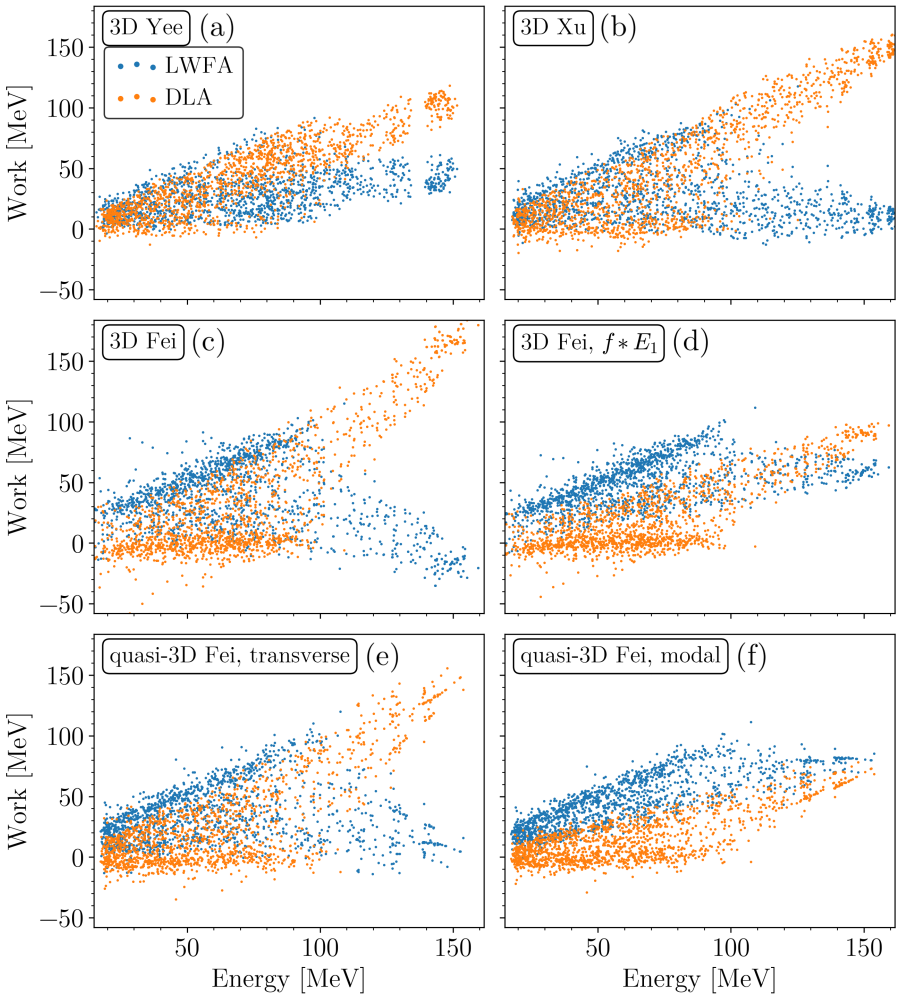

Figure 3.14: Work done by the LWFA and DLA mechanisms as a function of total electron energy for the three-dimensional and quasi-3D simulations using the various solvers. (d) The quantity $f * E_1$ is added to W_{DLA} and subtracted from W_{LWFA} to approximately take into account the energy loss of electrons from the longitudinal component of the laser. The work in (e) and (f) is computed using Eqs. (3.3) and (3.4), respectively, where (f) is the most accurate computation.

relation. The Xu solver exhibits the correct phase velocity for light in vacuum (but retains the spurious Lorentz force), and it provides a DLA energy gain that is greater than for the Yee solver. From (f) we see that for these simulation parameters, the most energetic electrons actually gain approximately equal energy from the LWFA and DLA processes. However, the LWFA energy gain appears to plateau as the dephasing limit is reached, whereas the DLA energy gain is expected to increase continuously (to a point) with propagation distance.

Finally, we wish to examine the dependence on the energy gain with respect to the electron's location in phasespace and real space. To do this, we show in Fig. 3.15 tracked electrons colored by their energy gain in the (a) angle-energy space and (b) real space. Data points for the tracked particles are overlayed on the phasespace density (grayscale) for all electrons; the tracked particles are the red and blue dots. In each pane, the middle group of tracked electrons is colored as $W_{\mathrm{DLA}} - W_{\mathrm{LWFA}}$, the bottom overlay shows $-W_{\mathrm{LWFA}}$ and the top overlay shows W_{DLA}. Thus the middle group shows the relative strength of the LWFA and DLA processes and the bottom (top) group shows the LWFA (DLA) process alone. From the top group in (a) we see that DLA gives more energy to electrons with a larger resulting angle away from the center, especially at very high energies. In contrast, the bottom group in (a) shows only increasing LWFA contribution with energy.

Perhaps the most distinct trend is visible in the top selection of electrons in (b); we observe one group close to the central axis that gains little energy from DLA, then another group somewhat sinusoidally distributed about the axis that gains significant energy from DLA. Note the width of this distribution in z, highlighting that electrons of various phases with the laser field can achieve betatron resonance (defined as $N\omega_\beta = \omega_0 - k_0 v_z$ for some integer N). Finally, those electrons injected first (larger z values) have had more time to interact with the laser and have gained more energy. We note that the overall shape of the accelerated electron bunch in the grayscale does not appear to be shaped sinusoidally in space like the electron bunch from the Yee solver shown in Fig. 3.13(a). This sinusoidal shape has been a common feature identified in numerous simulation papers as evidence for

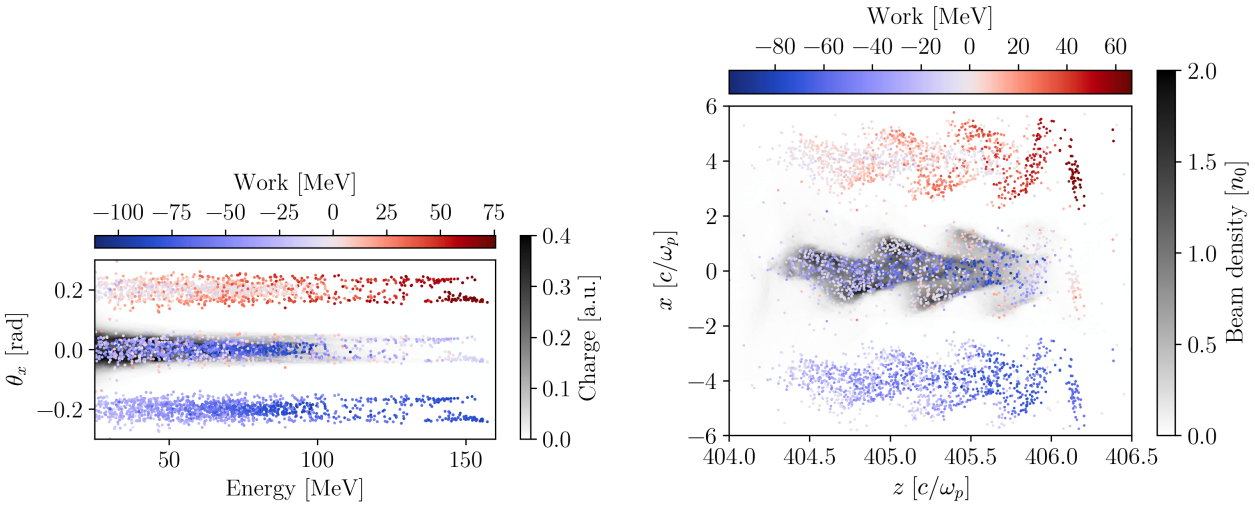

Figure 3.15: Snapshot of OSIRIS simulation data for the three-dimensional case using the Fei field solver before exiting the plasma. In both plots we show data for all electrons on a grayscale, with an overlay of individual tracked electrons in red/blue for (a) angle as a function of energy and (b) beam position. In each pane the middle overlay shows $W_{\mathrm{DLA}} - W_{\mathrm{LWFA}}$, the bottom overlay shows $-W_{\mathrm{LWFA}}$ and the top overlay shows W_{DLA}.

DLA [156, 1], but perhaps in many cases this shape is a result of errors in the phase velocity of the light wave. We explore this more fully in Sec. 3.3.2.

In conclusion, in this section we showed that using a finite-difference field solver that provides proper numerical dispersion and Lorentz force on particles in a laser gives substantially different results for an accelerated electron bunch experiencing DLA when compared to using the standard Yee solver. The transverse motion of the bunch no longer extends outside of the bubble radius, the transverse angle of the electron bunch is smaller and aligns more closely with experiment, and particles are observed to gain more energy from DLA due to the corrected phase velocity of the light. The characteristic sinusoidal distribution in space of an electron bunch experiencing DLA may not be accurate due to field solver errors, and the traditional longitudinal–transverse accounting for LWFA–DLA energy gain used in most PIC codes overestimates the DLA energy gain. Comparison of quasi-3D and full-3D simulations of an LWFA show that the quasi-3D algorithm can not only give consistent simulation results at a fraction of the computational cost, but also provide a more accurate calculation of DLA energy gain of electrons.

3.3 Examining betatron motion via numerical integration of the equations of motion

3.3.1 Equations of motion

Much work has been devoted to calculating the dynamics of electrons executing betatron motion within a laser field [68, 134], both in an ion channel [138] and in a wakefield [156]. Electrons undergoing betatron motion in a plasma wave wakefield can experience acceleration both from the longitudinal wakefield as well as from direct laser acceleration when the laser fields overlap the wakefield [134]. In Ref. [156], the equation of motion for such electrons is solved numerically and compared to results from a PIC simulation. However, the phase

velocity of light used in Ref. [156] is less than the speed of light due to dispersion errors in the PIC results. We therefore wish to numerically solve this equation of motion with the proper phase velocity, examining cases of DLA in both a wakefield and an ion channel. Laser fields can either be assumed to be those of a plane wave or those inside a cylindrical waveguide, where a longitudinal field associated with the laser is present in the latter case.

We first write out the equations of motion for an electron in a wakefield and plane-wave laser field. In this section only, we change our coordinate axes (switch x and z) from that used in the rest of this chapter. The laser propagates in the x-direction with wavenumber k_x and frequency ω_L, where the phase velocity of the laser is given as $v_p = \omega_L/k_x$. We introduce the coordinates $\xi \equiv x - v_p t$ and $\zeta \equiv x - v_g t$, where v_g is the group velocity of the laser. For the cases we are working with, we have that $v_g < c < v_p$, with the value of each taken from OSIRIS simulations. We assume the bubble to be spherical in shape with radius $R(x)$ (assuming $\frac{\partial R}{\partial x} \ll 1$), and we take the origin to be at the front of the bubble (back of the bubble initially located at $x = -2R$). For a plane-wave laser, we write the vector potential as $\boldsymbol{A}(x,t) = A_L \cos(k_x \xi)\hat{z}$, resulting in fields well approximated by $\boldsymbol{E}_L = -\frac{\partial \boldsymbol{A}}{\partial t} = -A_L \omega_L \sin(k_x \xi)\hat{z}$ and $\boldsymbol{B}_L = -\frac{\partial A_z}{\partial x}\hat{y} = A_L k_x \sin(k_x \xi)\hat{y}$. Here we have assumed that $\frac{\partial A_L}{\partial t} \ll A_L \omega_L$ and that A_L does not vary with ζ (as is assumed in Ref. [156]). The focusing and accelerating fields of the wake are of equal magnitude. We write the static longitudinal field, $\boldsymbol{E}_S$ as $E_{S,x} = \frac{\Gamma \tilde{n}(x)\omega_p^2 m_e}{2e}(x + R(x) - v_g t)$, where m_e and e are the mass and charge of an electron, respectively (e taken to be positive), ω_p is the plasma frequency for a plasma of density n_0 assuming complete blowout, $\tilde{n}(x) \equiv n(x)/n_0$, and Γ is a flag that is set to either zero or one. We make use of the $\tilde{n}(x)$ function to allow for a slow density ramp or modulation in the x-direction, and by setting Γ to zero the same equations can be used to model motion in an ion channel. The static transverse focusing field is $E_{S,z} = \frac{\tilde{n}(x)\omega_p^2 m_e}{2e} z$ for all z (to stop any electrons from escaping transversely).

To account for errors in the standard field solver on the Yee mesh used in most PIC codes, we recall that there is an error in the Lorentz force term for a particle moving in a laser

field in vacuum using standard PIC algorithms due to the electric and magnetic field being staggered in time from one another. Specifically, the force on a particle in the plane wave fields described above becomes

$$\boldsymbol{F} = q(\boldsymbol{E}_L + \boldsymbol{v} \times \boldsymbol{B}_L \cos\frac{\omega_L \Delta t}{2}), \tag{3.22}$$

where Δt is the time step used in the PIC simulation. Thus we introduce the factor τ_{stag}, which we multiply to all force terms containing $\boldsymbol{B}_L$ terms. When ignoring the time-stagger errors, we have $\tau_{\text{stag}} = 1$; to include these errors, we set $\tau_{\text{stag}} = \cos\frac{\omega_L \Delta t}{2}$ for some pre-determined value of Δt.

Using the fields and factors described above, we write the equations of motion for an electron as

$$\begin{aligned}
\frac{d}{dt}(\gamma m_e \dot{x}) &= -e(E_{S,x} - \tau_{\text{stag}} \dot{z} B_{L,y}) = -\frac{\Gamma \tilde{n}(x) \omega_p^2 m_e}{2}(x + R - v_g t) + eA_L k_x \tau_{\text{stag}} \dot{z} \sin(k_x \xi), \\
\frac{d}{dt}(\gamma m_e \dot{z}) &= -e(E_{S,z} + E_{L,z} + \tau_{\text{stag}} \dot{x} B_{L,y}) = -\frac{\tilde{n}(x) \omega_p^2 m_e}{2} z + eA_L k_x (v_p - \tau_{\text{stag}} \dot{x}) \sin(k_x \xi),
\end{aligned} \tag{3.23}$$

where a dot represents a time derivative. By differentiating the left-hand side of Eq. (3.23) and dividing by m_e and γ, we can write these equations as

$$\begin{aligned}
\ddot{x} + \frac{\dot{\gamma}}{\gamma}\dot{x} + \Gamma \tilde{n} \omega_\beta^2 x &= \Gamma \tilde{n} \omega_\beta^2 (v_g t - R) + \frac{a_0 c k_x \tau_{\text{stag}}}{\gamma} \dot{z} \sin(k_x \xi), \\
\ddot{z} + \frac{\dot{\gamma}}{\gamma}\dot{z} + \tilde{n} \omega_\beta^2 z &= \frac{a_0 c k_x}{\gamma}(v_p - \tau_{\text{stag}} \dot{x}) \sin(k_x \xi),
\end{aligned} \tag{3.24}$$

where $\omega_\beta = \omega_p / \sqrt{2\gamma}$ is the bounce frequency at which electrons naturally oscillate in the transverse direction, and $a_0 = \frac{eA_L}{m_e c}$ is the normalized vector potential.

We can more accurately describe the laser field by considering it to evolve inside of a cylindrical waveguide. Although this is inaccurate in the first bucket due to the rapidly changing bubble radius, it nonetheless gives a reasonable estimate of the longitudinal laser field. To do this we use the fields given in Ref. [134], namely,

$$\begin{aligned}
\boldsymbol{E}_L &= A_L \omega_L \frac{k_\perp}{k_x} J_1(k_\perp z) \cos(k_x \xi) \hat{x} - A_L \omega_L J_0(k_\perp z) \sin(k_x \xi) \hat{z}, \\
\boldsymbol{B}_L &= \frac{A_L \omega_L}{v_p} \left\{ J_0(k_\perp z) + \frac{k_\perp^2}{2k_x^2} \left[J_0(k_\perp z) - J_2(k_\perp z) \right] \right\} \sin(k_x \xi) \hat{y},
\end{aligned} \tag{3.25}$$

where $J_n(x)$ are Bessel functions of the first kind and order n, and $k_\perp = \frac{2.4}{R+1\, c/\omega_p}$. Solving then for the electron equations of motion yields

$$\begin{aligned}
\ddot{x} + \frac{\dot{\gamma}}{\gamma}\dot{x} + \Gamma\tilde{n}\omega_\beta^2 x =& \Gamma\tilde{n}\omega_\beta^2(v_g t - R) \\
&- \frac{a_0 c k_x}{\gamma}\Bigg\{ v_p \frac{k_\perp}{k_x} J_1(k_\perp z)\cos(k_x\xi) \\
&\qquad - \tau_{\text{stag}}\dot{z}\left[J_0(k_\perp z)\left(1 + \frac{k_\perp^2}{2k_x^2}\right) - \frac{k_\perp^2}{2k_x^2} J_2(k_\perp z)\right]\sin(k_x\xi)\Bigg\}, \\
\ddot{z} + \frac{\dot{\gamma}}{\gamma}\dot{z} + \tilde{n}\omega_\beta^2 z =& \frac{a_0 c k_x}{\gamma}\Bigg\{ J_0(k_\perp z)\left[v_p - \tau_{\text{stag}}\dot{x}\left(1 + \frac{k_\perp^2}{2k_x^2}\right)\right] \\
&\qquad + \tau_{\text{stag}}\dot{x}\frac{k_\perp^2}{2k_x^2} J_2(k_\perp z)\Bigg\}\sin(k_x\xi).
\end{aligned} \tag{3.26}$$

3.3.2 Collective electron behavior

We now seek to numerically integrate the partial differential equations in Eqs. (3.24) and (3.26) to determine the "correct" beam shape and angle of propagation, given the differing results from the PIC simulations. We perform this integration using a numerical solver for systems of ordinary differential equations[2] and follow a similar strategy to that described in Ref. [156]. Using the observed injection time from the PIC simulation as a guide, we inject particles into the bubble over a time of $200\,\omega_p^{-1}$. These particles are injected on-axis, with ζ chosen randomly from the interval $\zeta \in [-2R + 0.286\, c/\omega_p, -2R + 0.836\, c/\omega_p]$ and z from the interval $z \in [-r_b(\zeta), r_b(\zeta)]$, where $r_b(\zeta)$ is the bubble radius as a function of ζ. The initial particle energy is randomly selected from $\gamma \in [\gamma_g, 2\gamma_g]$, where $\gamma_g = (1 - (v_g/c)^2)^{-1/2}$. The angle of initial velocity to the x-axis is randomly chosen with the constraint that $v_x \geq v_g$. We show characteristics of the beam in Figs. 3.16–3.18 after a propagation time of $255\,\omega_p^{-1}$

[2]We use the scipy.integrate.solve_ivp solver in Python with an explicit Runge-Kutta method of order 5(4). The error is controlled assuming accuracy of the fourth-order method, but steps are taken using the fifth-order accurate formula (local extrapolation is done). A quartic interpolation polynomial is used for the dense output. A variable time step is used internally within the solver.

from when the first electron was injected. This is done for four different cases: in (a) and (b) we solve Eq. (3.24) for a plane-wave laser, and in (c) and (d) we solve Eq. (3.26) for a laser in a waveguide. For (a) and (c) we set $\tau_{\text{stag}} = 1$, but for (b) and (c) we set $\tau_{\text{stag}} = \cos\frac{\omega_L \Delta t}{2}$, where $\Delta t = 0.01875\,\omega_p^{-1}$. In all cases we use $v_p = 1.002565\,c$, $v_g = 0.993583\,c$, $k_x = 12.14\,\omega_p/c$ and $R = 3\,c/\omega_p$. For the plane-wave and waveguide cases, we have that $\omega_L = ck_x$ and $\omega_L = c\sqrt{k_x^2 + k_\perp^2}$, respectively. For both cases we use $a_0 = 34\,\omega_p/\omega_L$, giving 2.80066 and 2.79724 for the plane-wave and waveguide cases, respectively. These parameters are modeled to coincide with the three-dimensional OSIRIS simulations presented in Sec. 3.2.4.

The accelerated electron bunch is shown in real space at a time of $255\,\omega_p^{-1}$ in Fig. 3.16. We observe little difference between the plane-wave and waveguide laser models, but the time-stagger error shown in (b) and (d) increases the amplitude of transverse oscillations of the electrons. Some electrons even oscillate past the bubble radius, similar to the PIC simulations using the Yee solver pictured in Fig. 3.13(a). We also note the presence of a central, on-axis group of electrons in Fig. 3.16(c) that do not oscillate much with the laser fields. In (d), where the spurious Lorentz force is present, this central group of electrons exhibits large transverse displacements and is distributed sinusoidally in ζ, once again similar to the behavior observed using the Yee solver in Fig. 3.13(a). From these results we conclude that using a field solver that eliminates the spurious Lorentz force discussed in Sec. 3.2.2 is crucial in determining the spatial distribution of an electron bunch with overlapping laser fields. In addition, any errors in the electron motion will also affect the calculation of any betatron radiation produced by these oscillating electrons.

Next we examine the angle versus energy phasespace and real space distributions of accelerated electrons for the same four cases. For all panes in Figs. 3.17 and 3.18, the resulting distribution of electrons after a time of $255\,\omega_p^{-1}$ is shown in the middle of the pane. This middle group is colored as $W_{\text{DLA}} - W_{\text{LWFA}}$, or the work done by the DLA process minus the work done by the wakefield. Blue electrons thus indicate more energy received from the wake than from DLA, and vice-versa for red electrons. The top and bottom electron groups

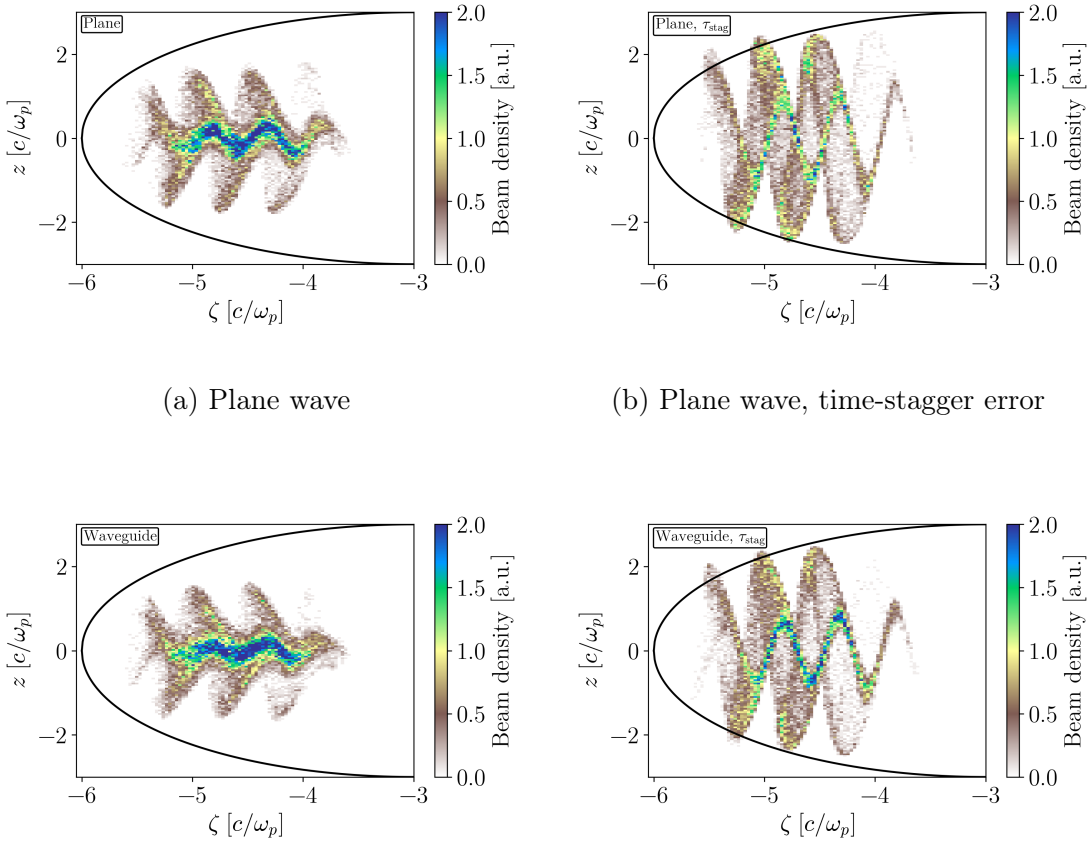

(a) Plane wave

(b) Plane wave, time-stagger error

(c) Waveguide

(d) Waveguide, time-stagger error

Figure 3.16: Resulting electron beam density for 30,000 test electrons injected and accelerated for a time of $255\,\omega_p^{-1}$. Using a waveguide laser model instead of a plane wave has little effect on the beam shape, while the time-stagger error enhances transverse oscillations (electrons even extend past the bubble radius).

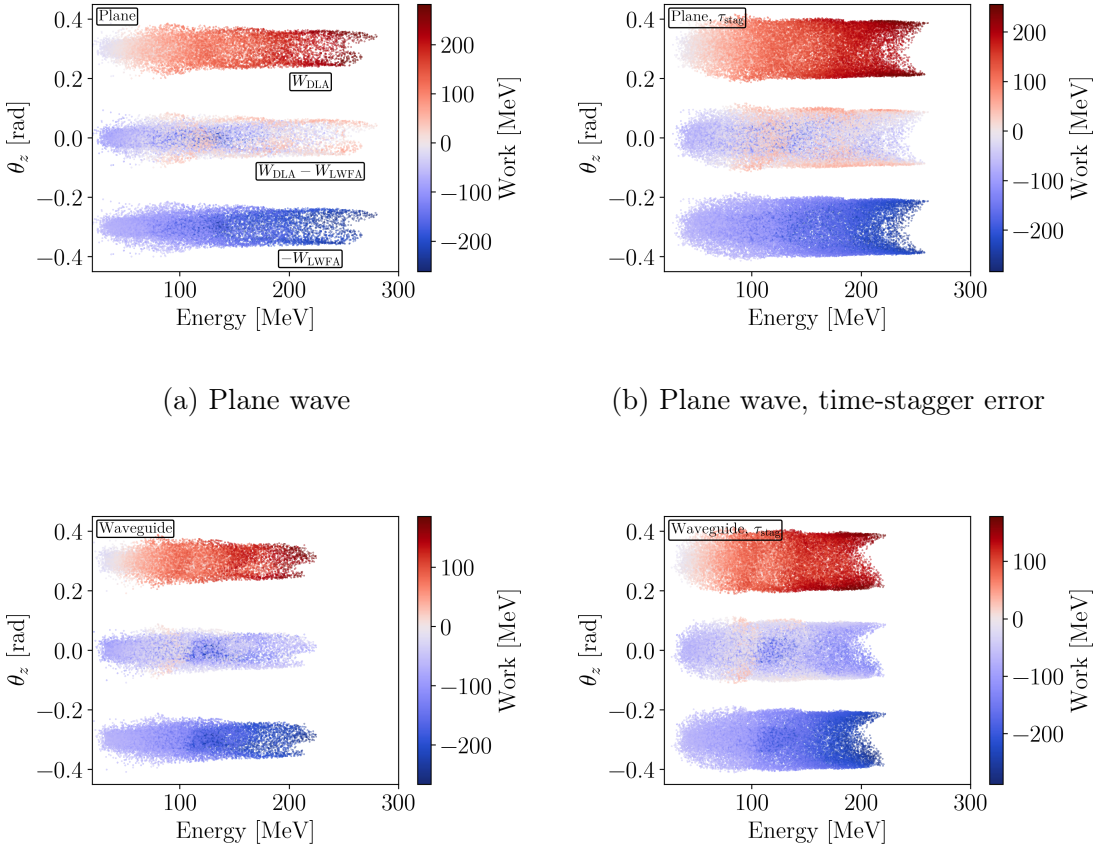

(a) Plane wave

(b) Plane wave, time-stagger error

(c) Waveguide

(d) Waveguide, time-stagger error

Figure 3.17: Transverse angle of propagation as a function of energy for 30,000 test electrons after a time of $255\,\omega_p^{-1}$. In all panes, the middle group is colored according to $W_{\mathrm{DLA}} - W_{\mathrm{LWFA}}$, or the work done by the DLA process minus the work done by the wakefield. The top and bottom groups of electrons are offset for visualization and colored as W_{DLA} and $-W_{\mathrm{LWFA}}$, respectively. Once again, the time-stagger error increases transverse momentum. Here we also see that the plane wave laser contributes increased DLA energy due to the larger field strength near the bubble radius.

shown in each pane are vertically offset for visualization purposes; they are colored as W_{DLA} and $-W_{\mathrm{LWFA}}$, respectively.

Figure 3.17 shows the transverse angle of propagation as a function of electron energy. The forking at high energies is characteristic of the DLA process, and the range of θ_z is larger when the time-stagger error is present, similar to the OSIRIS simulation results shown in Figs. 3.12 and 3.15. We also note the presence of two distinct electron groups for each case shown in Fig. 3.17: one group gains little energy from DLA and is located close to the central axis, and the other group gains significant energy from DLA and clusters at larger values of θ_z. This is because electrons originating at large initial values of x execute betatron oscillations with a larger average transverse momentum. In (a) the first group has an energy of approximately 25–150 MeV and the second group has energy up to 280 MeV. The results shown here highlight that a spurious Lorentz force on resonant electrons erroneously increases the amplitude of their transverse oscillations.

We also observe two distinct electron groups in Fig. 3.18: the wakefield-accelerated electrons (blue) are only slightly modulated away from the axis, while the DLA-accelerated electrons (light red) are much farther displaced. In Fig. 3.18(b) and (d), the wakefield-accelerated electrons are displaced in near-sinusoidal pattern away from the axis as a result of the time-stagger error. Many important beam quantities are inaccurate when a solver is used that exhibits a time-stagger error, as is the case for most PIC codes. The spatial distribution of electrons executing betatron motion in (c) closely resembles the shape of the electron bunch from the OSIRIS simulation shown in Figs. 3.13(c) and 3.15(b), where the Fei solver is used. From this comparison, we conclude that the use of the Fei solver [37], which corrects the numerical dispersion relation and removes a spurious Lorentz force on particles, indeed generates improved results that more accurately resemble the desired physics.

Finally, in Fig. 3.19 we plot the spatial distribution of the test electrons again late in time, but here colored according to their initial values of (a) x and (b) v_x. We see that electrons are more likely to be located in the DLA group (executing large transverse oscillations) given

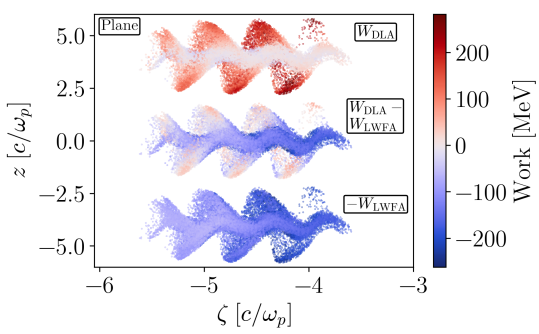

(a) Plane wave

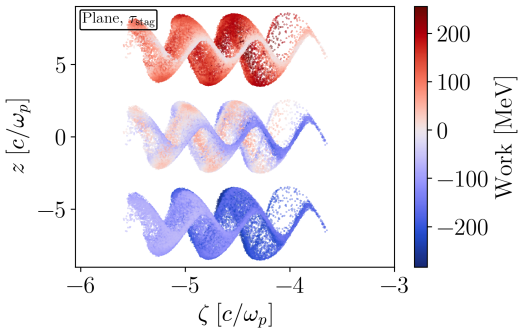

(b) Plane wave, time-stagger error

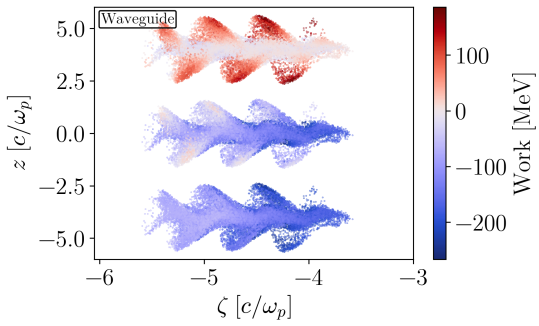

(c) Waveguide

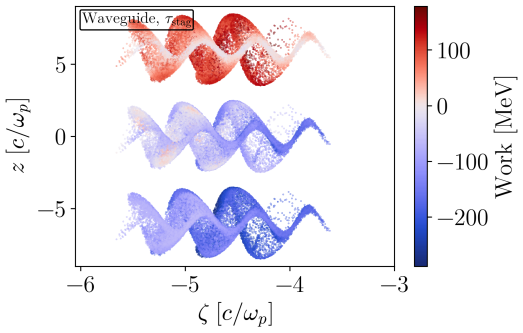

(d) Waveguide, time-stagger error

Figure 3.18: Spatial snapshot of 30,000 test electrons after a time of $255\,\omega_p^{-1}$, where the groups in each pane are colored the same as in Fig. 3.17. For all cases, we observe a central line of electrons nearly unaffected by DLA; however, this bunch is displaced from the axis when the time-stagger error is present. In all cases, DLA most strongly accelerates the electrons with large transverse displacement.

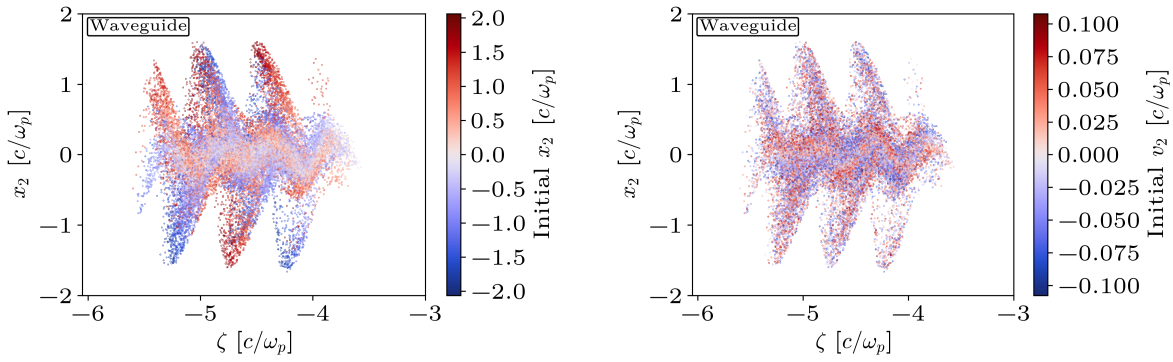

(a) Plane wave

(b) Plane wave, time-stagger error

Figure 3.19: Spatial snapshot of 30,000 test electrons after a time of $255\,\omega_p^{-1}$ for the waveguide case with no time-stagger error. The groups in each pane are colored to show the initial values of (a) x and (b) v_x for each particle.

an initially large transverse position, but this large transverse position does not guarantee a particle to be located in the DLA group. In addition, no significant trend is observed with the initial transverse velocity. (The initial phase relative to the laser was also plotted, which likewise showed no significant trend.) Thus, injection schemes where electrons are injected some distance away from the laser axis can result in a greater population of electrons in betatron resonance.

3.3.3 Single-particle betatron motion

To further understand the acceleration process of electrons from Sec. 3.3.2, we integrate the equations of motion in Eq. (3.26) without the time-stagger error for 1,000 electrons. We also integrate out to longer times to more fully explore the dynamics. In Fig. 3.20, we show two representative particles, (a) one dominated by the wakefield and (b) another receiving comparable energy from both the LWFA and DLA mechanisms for the plane-wave case. In the top pane we see the transverse velocity, v_z, along with (minus) the laser electric field experienced by each particle. These quantities are observed to be in phase while the electron

in (b) gains energy from DLA. Note that in vacuum, where transverse canonical momentum would be conserved, they would be $\pi/2$ out of phase. In an ion channel or wakefield, there is an additional electric (and magnetic) field that breaks translational invariance in the transverse direction. In the second pane, we plot the betatron frequency, $\omega_\beta = \omega_p/\sqrt{2\gamma}$, and the Doppler-shifted laser frequency, $\omega_D = (1 - v_x/v_p)\,\omega_L$, for each particle. At the betatron resonance we have $N\omega_\beta \approx \omega_D$; the quantities ω_β and ω_D roughly overlap during positive DLA energy gain in (b) at the first harmonic ($N = 1$). Finally, in the bottom pane we observe the total energy of the particle, along with the cumulative contributions from the LWFA and DLA mechanisms. In (a) the electron never establishes prolonged resonance during the majority of its time in the bubble, and the γ curve closely follows the LWFA energy gain, peaking at 130 MeV. However, in (b) the energy gains from LWFA and DLA are of similar magnitude, resulting in a maximum total energy of 290 MeV. For these parameters, the DLA mechanism is able to produce electrons with double the total energy compared to the wakefield mechanism alone.

Next, we investigate the maximum energy gain accessible to an electron in betatron resonance from DLA. There are many interesting research directions to pursue. For example, for an electron in an ion channel (where it won't dephase and slow down from the wake), what is the maximum energy an electron can gain from the betatron resonance? Can a density gradient or counter-propagating laser be introduced to maintain the betatron resonance to attain higher energy? Can smooth transitions between successive harmonics occur over the energy gain, or can a stochastic process accomplish the transition?

As an electron is accelerated by the laser fields while in resonance ($N\omega_\beta \approx \omega_D$), the γ factor for the particle increases, causing ω_β to decrease. However, the forward velocity v_x of the particle remains largely unchanged (aside from the oscillation due to the particle moving transversely across the channel width). Thus if ω_β decreases too much, the particle will fall out of resonance and begin to lose energy. This is observed in the top and bottom panes of Fig. 3.20(b); we see that the electron transitions into anti-resonance at a time of

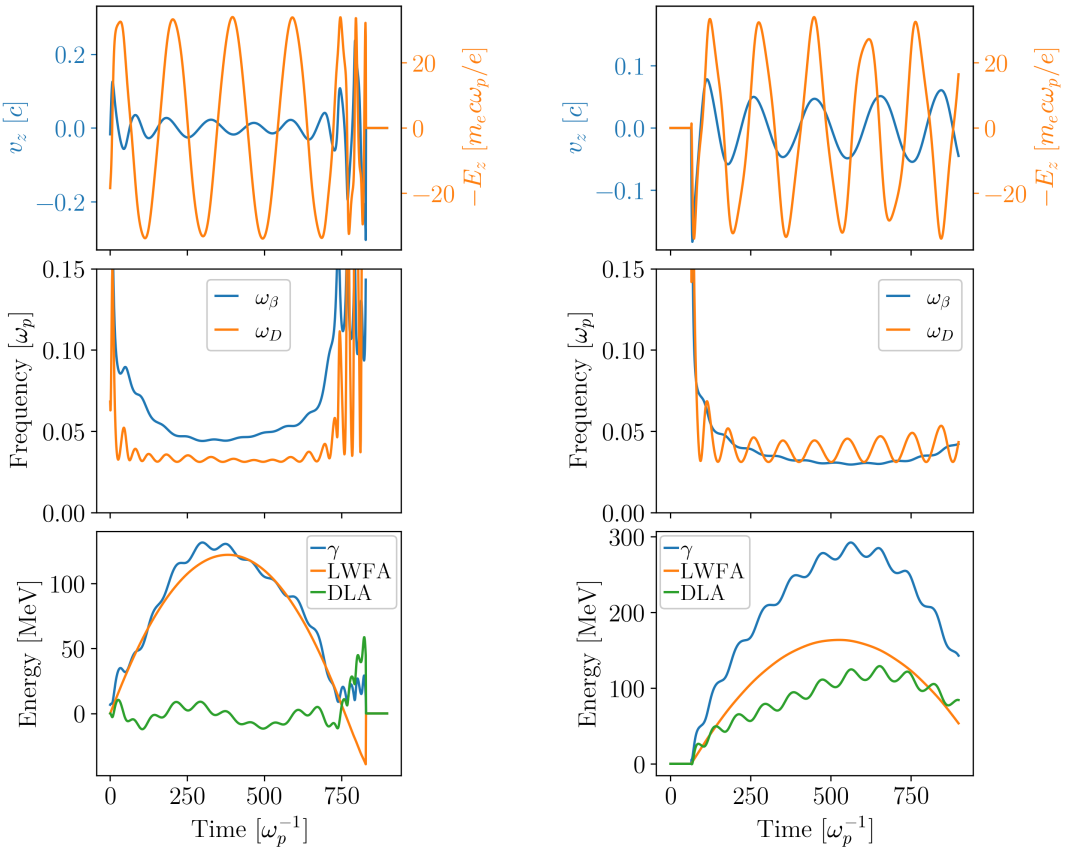

(a) Wake-dominated particle

(b) Wake- and DLA-dominated particle

Figure 3.20: Trajectories for two electrons, (a) one dominated by the wakefield and (b) another receiving roughly equal energy from the LWFA and DLA mechanisms. Both particles eventually dephase and lose energy from the wakefield, but we see signatures of betatron resonance for the particle in (b).

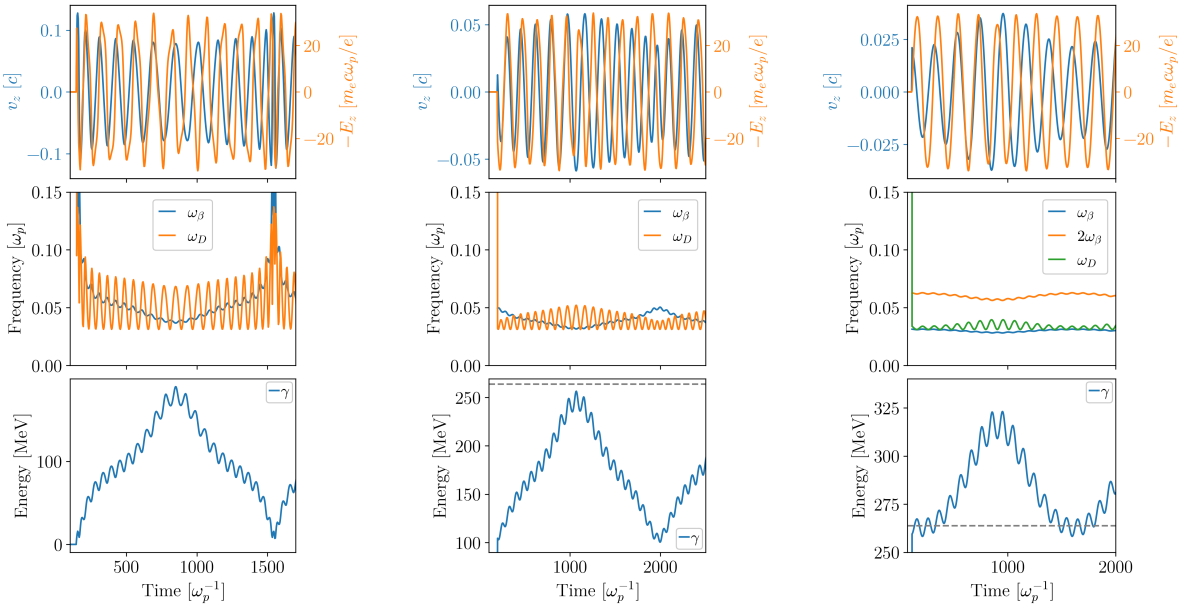

(a) $\gamma_i = \gamma_g = 8.84$ (b) $\gamma_i = 200$ (c) $\gamma_i = 500$

Figure 3.21: Electron trajectories within an ion channel (no longitudinal wakefield) for three different initial energies. In each case, the electron cyclically gains and loses energy to the laser fields as it moves in and out of resonance. The gray dashed line in the bottom pane of (b) and (c) indicates the $N = 1$ betatron resonance condition for a particle with no transverse velocity, illustrating that for this case the overall energy gain does not extend across multiple orders of resonance.

around $700\,\omega_p^{-1}$, characterized by the v_z and $-E_z$ curves being opposite and the DLA energy decreasing. In fact, Ref. [138] shows that the trajectories of resonant particles either oscillate or rotate in the θ-p_x phasespace, where θ is the phase difference between the laser electric field and the transverse velocity. Thus in the presence of only a laser electric field inside an ion channel, there will always be a maximum energy gain for electrons executing betatron motion.

To further investigate this process, we again integrate the equations of motion in Eq. (3.26) without the time-stagger error for 1,000 electrons, but this time we remove the longitudinal

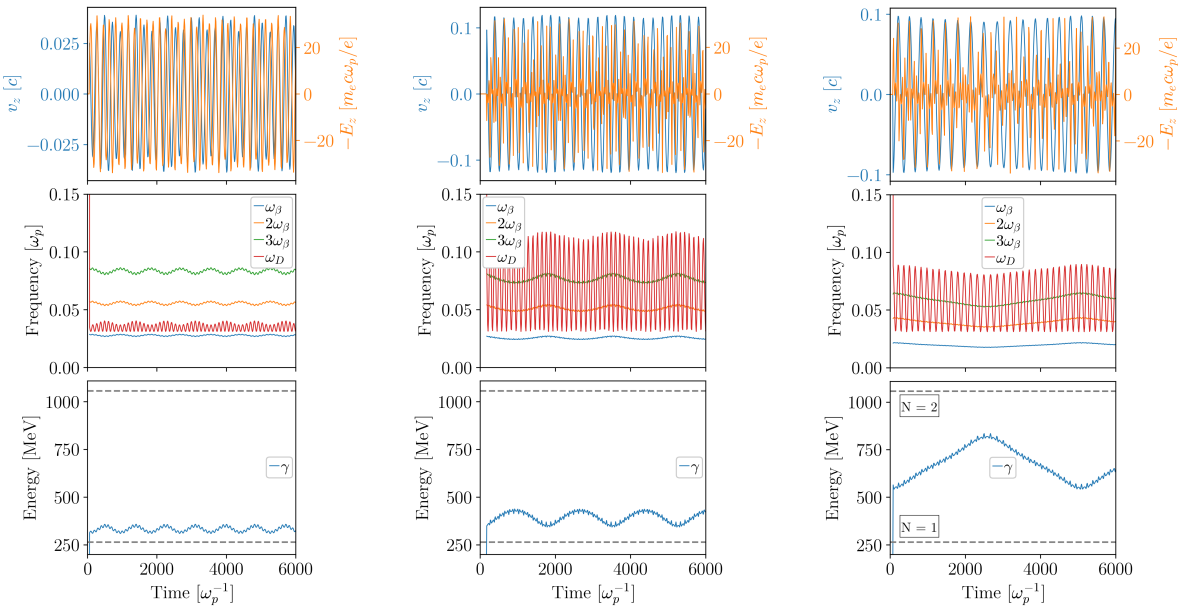

(a) $\gamma_i = 600$ (b) $\gamma_i = 700$ (c) $\gamma_i = 1100$

Figure 3.22: Electron trajectories within an ion channel (no longitudinal wakefield) for three different initial energies in-between the first- and second-order resonances. In each case, the electron cyclically gains and loses energy to the laser fields as it moves in and out of resonance. It appears that the electron is unlikely to spontaneously jump from the first- to the second-order resonance.

wakefield force (set $\Gamma = 0$). We initialize particles with energy in the range $\gamma \in [\gamma_i, \gamma_i + \gamma_g]$, where γ_i is set to some value for each integration and $\gamma_g = 8.84$ for the group velocity given above. To interpret the appropriate value for γ_i, we first find the values of γ that correspond to various resonance orders. For a particle with zero transverse velocity, our resonance condition $N\omega_\beta = \omega_D$ becomes $N\omega_p/\sqrt{2\gamma} = (1 - \sqrt{1-\gamma^{-2}}/v_p)\omega_L$. Solving this equation numerically yields $\gamma_N = 516.154$, 2067.45 and 4652.11 for $N = 1$, 2 and 3, respectively. In Fig. 3.21, we plot single-particle trajectories for three different initial energies, $\gamma_i = (8.84, 200, 500)$, all smaller than γ_1. The resonance condition of $\omega_\beta = \omega_D$ is observed to be loosely satisfied in each case.

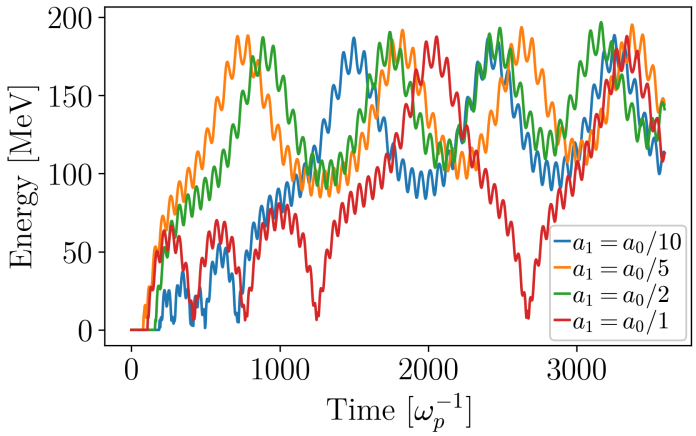

Figure 3.23: Electron kinetic energies as a function of time with a counter-propagating laser of amplitude a_1. The electron trajectories are no longer strictly periodic, but there is little increase in maximum energy.

We also show in Fig. 3.22 the trajectories of electrons with starting energies between the first- and second-order resonances, $\gamma_i = (600, 700, 1100)$. We see that the first electron in (a) remains at the first-order resonance during its entire propagation, but that its energy does slightly increase. However, although the electron in (b) begins with nearly the same energy as the maximum attained by the electron in (a), it exhibits a second-order resonance. The oscillatory behavior of the energy is at a lower frequency—and lower still in (c)—and the energy gain is larger. From these results it is unclear if particles are able to naturally transition to higher-order resonances on their own. This is an area for future research.

3.3.4 Impact of a counter-propagating laser

To introduce some possible stochasticity into the electron motion, we now include a counter-propagating laser of smaller amplitude into the equations of motion [157, 158, 159]. For the finite-size laser, the fields are the same as in Eq. (3.25), except ξ is replaced with $\xi^* = x + v_p t$ and the magnetic field is negated (the phase is shifted by π). We then leave the amplitude

of the main laser pulse at $a_0 = 2.79724$ and vary the magnitude of the counter-propagating laser pulse, a_1, between 10% and 100% of the value of a_0. The kinetic energy of the most energetic electrons for each test case are plotted together as a function of time in Fig. 3.23. We see that the electron trajectories are no longer strictly periodic as before, but that the energies still oscillate with significant variation. For the cases with $a_1 \leq a_0/2$, we see that the electrons remain at energies larger than $\sim 100\,\mathrm{MeV}$. The use of such a counter-propagating laser could be advantageous for maintaining high electron energies over large distances in an ion channel. This is also an area for future work. However, it is possible that noise inherent in an experiment is enough to inherently disrupt the periodicity of the electron motion.

3.4 Simulating an electron probe beam diagnostic

In order to experimentally study in detail the dynamics of particles in combined wakefields and laser fields, it is important that novel diagnostics be developed. There is thus interest in using an energetic (tens of MeV) proton or electron bunch perpendicularly incident on an LWFA- or SM-LWFA-generated wakefield to image the real-time, in situ fields inside a plasma. The particle bunch needs to be of short duration so that the displacement due to the wake and laser fields is coherent. In a recent experiment performed at the Accelerator Test Facility (ATF) at Brookhaven National Laboratory, an energetic electron beam illuminated an SM-LWFA-generated wakefield to probe the instantaneous characteristics of the wakefield. A 2-J and 2-ps CO_2 laser was used as the driver; this longer wavelength laser produces a longer wavelength wake, which is in turn easier to probe. The laser propagated through hydrogen gas, and an electron beam probe with an energy of 55 MeV was incident normal to the propagation of the wakefield and parallel to the laser polarization direction. In this setup we consider the laser to be polarized in the y-direction and to propagate in the z-direction. The electron beam incident on the plasma had a duration of 150 fs, a charge of 25 pC, and propagated in the y-direction with finite transverse sizes of σ_x and σ_z. A schematic of the

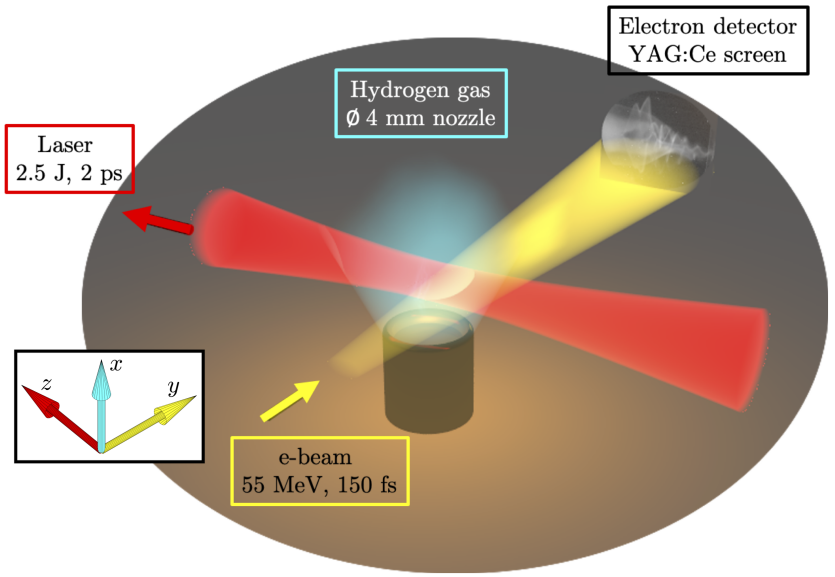

- **Electron beam:**
  - $E_b = 55$ MeV;
  - $\tau_e = 150$-250 fs;
  - $Q = 20$-30 pC.
- **CO_2 laser:**
  - $E_L = 1$-5 J;
  - $\tau_L = 2$ ps.

Figure 3.24: Schematic showing the hydrogen gas, CO_2 laser and electron probe beam used in the experiment performed at the ATF. The probe beam was delayed in time over various shots to construct multiple images of the laser propagating through the plasma. The distance from the plasma to the screen was 50 cm.

experiment is shown in Fig. 3.24, where the electron beam can be seen to form an image at a screen 50 cm away from the plasma. Radiographs from the electron beam were obtained over many snapshots by temporally delaying the beam with respect to the laser, showing the laser propagating through the plasma and evidence of the plasma evolution behind the laser. The snapshots at different delays then provide a "movie" of the wakefield evolution, provided that the SM-LWFA processes are consistent between shots.

The electron beam was passed through several focusing optics to create the desired spot size and divergence in the horizontal (z) and vertical (x) directions, with the beam

parameters being substantially different in the two directions. The beam dynamics in the ATF beamline 1 were simulated using `elegant` [160], and the horizontal and vertical beam envelopes throughout the beamline are shown in Fig. 3.25. The initial beam distribution was generated at the exit of the ATF RF linac with a normalized emittance of $\epsilon_n = [2.62,\ 2.89]\,\mu$m, where we express quantities here as $[*_x,\ *_z]$, i.e., separately for the x- and z-directions. Blue and red dots are the RMS beam sizes measured at two beam profile monitors located in the experimental chamber at the locations of the plasma (first set) and image screen (second set). The beam is much wider in the z-direction so as to overlap many plasma periods. The beam focus in each direction is seen to be located before the plasma, nearly the same distance away as the screen is from the plasma. The measured beam sizes give a beam magnification of $M_{\mathrm{exp}} = [2.0,\ 2.7]$ from the plasma to the screen in the vertical and horizontal directions, respectively. The simulated beam had a spot size at the screen of $\sigma(y_{\mathrm{scr}}) = [0.738,\ 3.582]\,\mathrm{mm}$ and a magnification of $M_{\mathrm{sim}} = [1.65,\ 1.58]$ between plasma and screen. The magnification was also measured using a third method: by delaying the probe in relation to the laser, a series of radiographs was generated showing the laser moving through the plasma. Knowing the synchronization delay and measuring the motion of certain features in the radiograph (e.g., the front of the first bucket) over time, this motion was compared to the speed of light to determine that the images had a magnification factor of 3–4 in the horizontal direction. Thus, there remains some uncertainty regarding the exact magnification of the beam between these three methods.

In this section, we seek to simulate the electron probe beam displacement from the plasma wakefield associated with parameters for the experiment to better interpret the radiographs. Since an accurate simulation of this problem relies on a three-dimensional geometry, we use the quasi-3D geometry of OSIRIS to simulate the original SM-LWFA problem. We then use the fields obtained from the simulation to propagate many test particles transversely across the simulation box. This propagation is done in post-processing due to the complicated nature of such a beam in the quasi-3D geometry. After "pushing" the particles across the

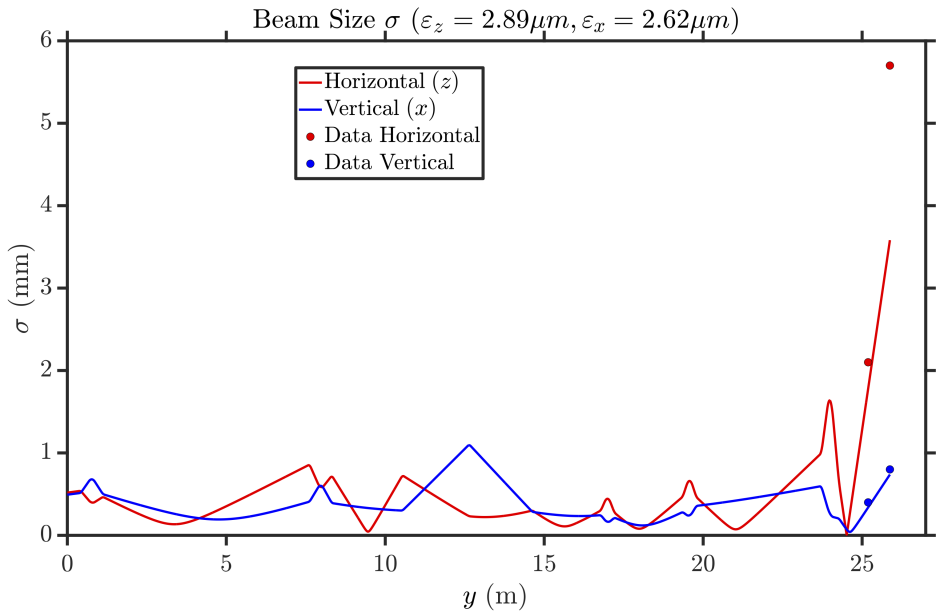

Figure 3.25: Simulated probe beam width in the horizontal and vertical directions as a function of propagation distance. Dots correspond to the experimental beam sizes taken in the absence of the plasma or screen (the first and second sets of dots correspond to the location of the plasma and screen, respectively). From the experimental (simulation) values, magnification is about 2.7 (1.6) in the horizontal direction.

simulation box, their straight-line trajectories are calculated to form an image at a screen some distance from the plasma. These images are then compared to the experimental results. Although the physics involved in this endeavor are not directly related to DLA, we include the results in this chapter due to this technique's potential to provide direct signatures for the wakefields. However, as we show, interpreting the images is complicated. The premise of the technique is that the probe electrons receive an impulse from the transient-time wakefields they encounter. They then free stream (self-fields are negligible since they are relativistic and the beam is not tightly focused) to the screen. If the screen is set before a focus for the perturbed beam, then the image can be more easily interpreted. However, more care must be undertaken when interpreting images from a screen placed beyond the beam focus.

Several nuances are involved when correctly modeling the electron probe beam. First, if electrons are initialized on a regular grid, this regular spacing can lead to aliasing of the field structures on the radiograph. We found that results were much better when the beam was initialized randomly. Second, when propagating the probe electrons across the simulation grid, we usually only had the field data at a single time step. Getting the data for all of the time steps during which the electrons cross the grid is too expensive (and unnecessary for temporally short probe beams). In most cases, we therefore assume the field structures to be drifting forward at some group velocity near the speed of light without any change to the envelope. This quasi-static approximation includes some error for a 150-fs electron beam, but we keep with this assumption as we develop a path towards modeling the process. Finally, we treat the fields to be zero outside of some radial simulation distance to ignore edge fields at the plasma–vacuum boundary as well as some large, spurious simulation boundary fields.

We want to initialize the electron beam at the focus, since only at that location are both position and momentum random and independent of one another, i.e., the distribution function is separable at the focus. Once again, we consider the laser to be polarized in the y-direction and to be propagating in the z-direction. We then let the electron beam propagate (free stream) in the y-direction with finite transverse sizes of σ_x and σ_z towards the plasma

edge. Without loss of generality, after deriving the necessary equations we can easily switch x and y for the probe beam so that the electrons propagate perpendicular to the polarization direction of the laser. We define $y = 0$ to be the focal point of the beam, with the plasma interaction plane (center of the wake) to be located at y_{IP} and the screen to be located at $y_{\mathrm{scr}} = y_{\mathrm{IP}} + L_{\mathrm{scr}}$, where L_{scr} is the distance between the plasma and the screen; we assume that L_{scr} is always known.

We have that the normalized emittance of the beam in a particular direction is $\epsilon_n = \sqrt{\langle x^2 \rangle \langle p_x^2 \rangle - \langle x p_x \rangle^2}$, where $p_x = \gamma x' = \gamma \frac{dx}{dy}$. Note that prime indicates derivative with respect to the propagation direction, $y = ct$. The normalized emittance is not conserved when a beam with energy spread propagates in vacuum. However, the geometrical emittance, $\epsilon = \epsilon_n / \gamma$, is always conserved in vacuum, and it defines the diffraction of a beam. Henceforth, we will work with the geometrical emittance, where $\beta = \frac{\sigma_x^2}{\epsilon}$. The equation of motion for a beam evolving in free space is [161]

$$\sigma_x'' = \frac{\epsilon^2}{\sigma_x^3}, \tag{3.27}$$

which has the solution

$$\sigma_x(y) = \sqrt{\frac{\epsilon^2 y^2}{\sigma_x^{*2}} + \sigma_x^{*2}} = \sigma_x^* \sqrt{1 + \frac{y^2}{\beta^{*2}}}, \tag{3.28}$$

where a superscript * indicates quantities defined at the focus. Note the analog between ϵ and β^* with λ and Rayleigh length, respectively, when comparing particle beam and laser propagation.

If we know the parameters y_{IP}, σ_x^* and β^*, i.e., we know the location of the focus and all of the beam parameters, then we have that the magnification of the beam from the plasma interaction plane to the screen is

$$M = \frac{\sigma_x(y_{\mathrm{scr}})}{\sigma_x(y_{\mathrm{IP}})} = \frac{\beta^{*2} + (y_{\mathrm{IP}} + L_{\mathrm{scr}})^2}{\beta^{*2} + y_{\mathrm{IP}}^2}. \tag{3.29}$$

If we instead know the beam emittance, spot size at the screen and magnification, we need to determine σ_x^* and y_{IP} to properly initialize the beam. This means we need to solve the

system of equations

$$\begin{aligned} \sigma_x(y_{\rm scr})^2 &= \frac{\epsilon^2(y_{\rm IP}+L_{\rm scr})^2}{\sigma_x^{*2}} + \sigma_x^{*2}, \\ M^2 &= \frac{\sigma_x^{*4}+\epsilon^2(y_{\rm IP}+L_{\rm scr})^2}{\sigma_x^{*4}+\epsilon^2 y_{\rm IP}^2} \end{aligned} \tag{3.30}$$

for σ_x^* and $y_{\rm IP}$, yielding

$$\begin{aligned} \sigma_x^* &= \frac{\epsilon L_{\rm scr} M}{\sqrt{(1+M^2)\sigma_x(y_{\rm scr})^2 - 2M\sqrt{\sigma_x(y_{\rm scr})^4 - \epsilon^2 L_{\rm scr}^2 M^2}}}, \\ y_{\rm IP} &= \frac{L_{\rm scr}(M^2-1)\sigma_x(y_{\rm scr})^2\left[\sigma_x(y_{\rm scr})^2 + M\sqrt{\sigma_x(y_{\rm scr})^4 - \epsilon^2 L_{\rm scr}^2 M^2}\right] - 2\epsilon^2 L_{\rm scr}^3 M^4}{4\epsilon^2 L_{\rm scr}^2 M^4 + (M^2-1)^2\sigma_x(y_{\rm scr})^4}. \end{aligned} \tag{3.31}$$

Using Eq. (3.31) we can properly initialize the particle beam. Because the beam is moving in the y-direction at near the speed of light, we can have separate locations of focus, spot sizes and emittances in the x and z directions.

Although the beam has a finite duration in time, we wish to push all particles through the plasma and laser fields at once, i.e., create one PIC loop that processes all the particles in the beam together. We accomplish this by initializing the particle beam with distributions in x, z and time, we then propagate all particles through vacuum to the edge of the plasma in y using straight-line trajectories.

We next assume that the wakefields are quasi-static; this includes all components of both the electric and magnetic fields, along with modes zero and one in the quasi-3D geometry. Given a field component from the simulation, $W(x,y,z;t_0)$, for a particular time t_0 from the OSIRIS simulation, we can assume that at a later time $t = t_0 + \Delta t$, the field is described as $W(x,y,z;t) = W(x,y,z-v_g\Delta t;t_0)$. Therefore, for beam particles that trail the head of the beam by a time $t_{\rm trail}$, they are simply started at the same t_0 but shifted in z by an amount $-v_g t_{\rm trail}$. In this way, all beam particles can be assumed to start at the same time and be pushed through the fields together via the Boris push. For any time step when a particle resides outside of the specified field domain, a straight-line trajectory is used. Finally, once all particles have passed through the domain, we use straight-line trajectories again to produce the image at the screen.

We simulated the ATF experiment using quasi-3D OSIRIS for a hydrogen plasma of density $n_0 = 10^{17}\,\mathrm{cm}^{-3}$ and a laser with wavelength $10.6\,\mu$m, peak amplitude $a_0 = 1.9$, beam waist $30\,\mu$m, and FWHM of intensity 2 ps polarized in the y-direction. We used mobile hydrogen ions with a Lorentzian for the longitudinal plasma profile of the form

$$n(z) = n_0 \left(-0.00132155 + \frac{1}{1 + 20\,|1.5 - z[\mathrm{mm}]|^{8.95735}} \right), \tag{3.32}$$

which approximated the experimentally measured density profile. The density was set to zero to the left and right of $z = 0$ and $3\,\mu$m, respectively, and the density profile has a nearly flat distribution over the middle $1\,\mu$m. In the transverse direction, the plasma tapered linearly from full density to zero from 300 to $400\,\mu$m, and the simulation boundary was placed in the transverse direction at $620\,\mu$m. The simulation was run for 43 ps to capture the evolution of the plasma after the laser passed through (no moving window), and had $c/\omega_p = 16.82\,\mu$m, $\Delta x_1 = 0.0105\,c/\omega_p$ and $\Delta x_2 = 0.0315\,c/\omega_p$, with 35,453 and 1,175 cells in x_1 and x_2, respectively.

The electron probe beam was simulated to pass through the fields at each data dump in post-processing. Fields were assumed to be constant in time and drifting at the group velocity $v_g = 0.9958959\,c$ obtained from the OSIRIS simulation data (little difference is observed if a group velocity of c is used). The electron probe beam was sent through traveling in the y-direction, parallel to the laser polarization. We show in Figs. 3.26–3.30 (a) the simulated electron radiograph at a screen 50 cm away from the plasma interaction plane, (b) the electron density [and (d) ion density] of the OSIRIS simulation, (c) a shot from the experimental radiograph roughly corresponding to the same time and features from the simulated radiograph, and (e) the laser electric field from the OSIRIS simulation. For better comparison between the plots, all axes are displayed in units of what is measured at the screen (except for the laser field); this means that the electron and ion densities are magnified by factors of 1.65 and 1.58 in the vertical and horizontal directions, respectively.

In Fig. 3.26, we can just barely see the displacement of plasma electrons on the color scale in (b). These plasma electrons are completely displaced due to the ponderomotive force

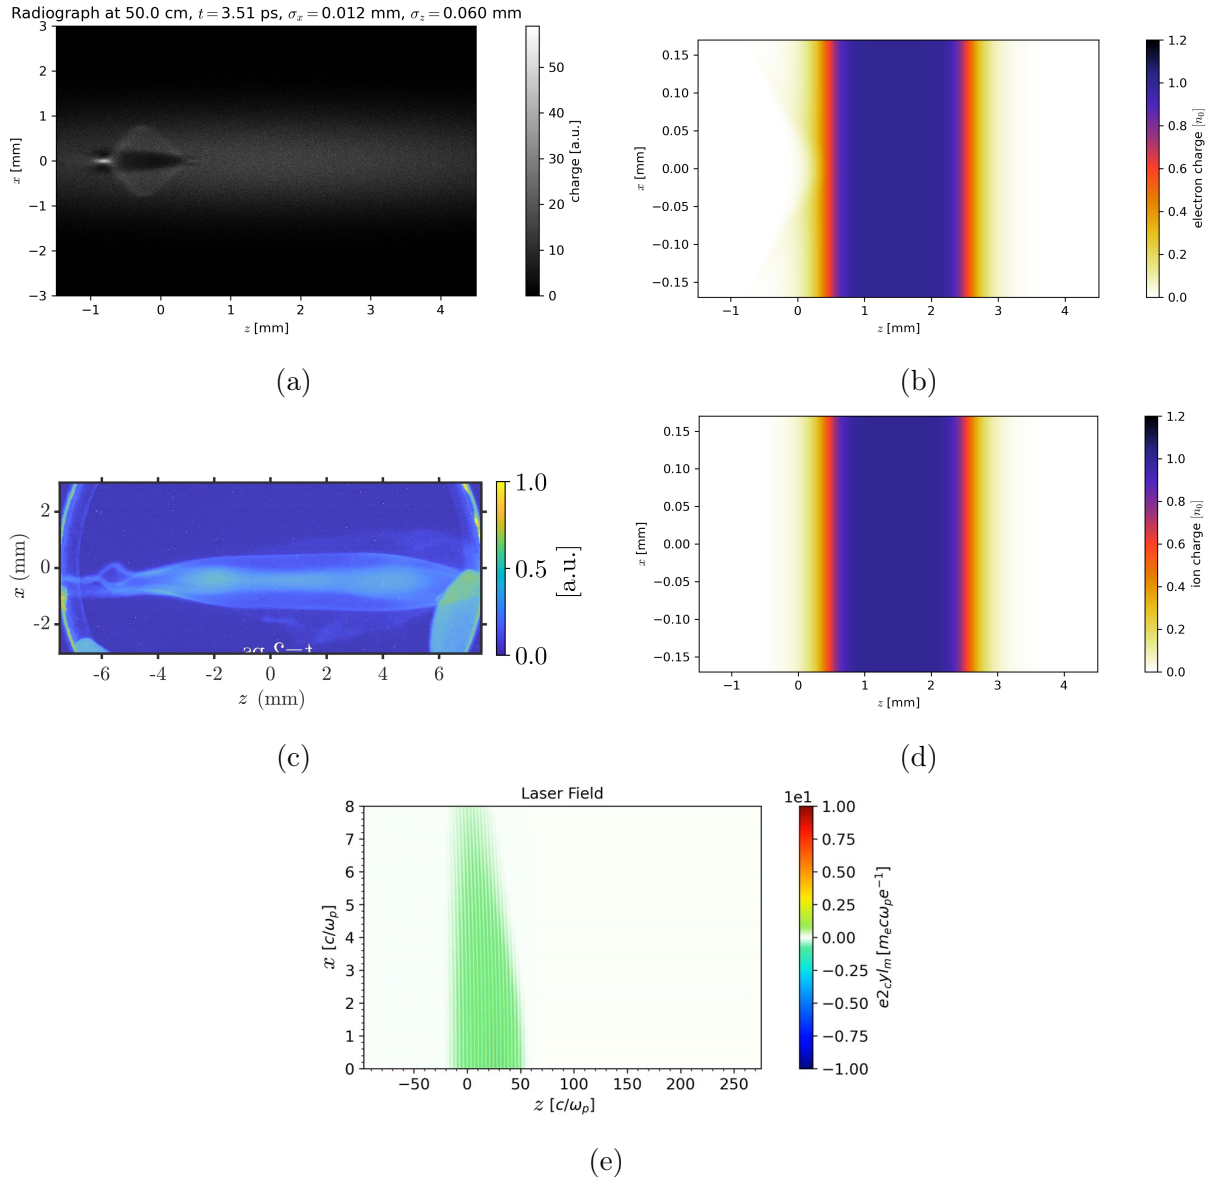

Figure 3.26: (a) Simulated electron radiograph produced on a screen. (b) Electron density just as the laser field enters the density upramp. (c) Experimental radiograph that resembles this time in the simulation. (d) Ion density, which has not yet been displaced. (e) Laser field of the quasi-3D OSIRIS simulation shown focusing to the axis.

of the laser, but the ions in (d) have not yet begun to move. Since the density is rapidly changing in the upramp, the electrons in the probe beam will experience a focusing force that varies in z due to the stationary ions. In (a), the radiograph shows that the electrons between about -0.5 and 0.5 mm have passed through their focus, while the electrons near -0.9 mm are focused at the screen. The varying space charge of the ions acts as a variable-strength focusing lens on the probe electrons. This facet of the radiography process highlights the difficulty of extracting plasma features from a single radiograph located a set distance from the plasma: any of the electrons at the screen may be imaged before, at, or after the focus of their trajectories. In future experiments, the distance of the image screen from the plasma will be varied to better interpret the data. However, we do observe a small group of electrons at around -6 mm in the experimental image in (c) that are displaced from the main probe beam, similar to the simulated probe.

Figure 3.27 shows a snapshot of when the laser has just approached the middle of the plasma. The laser is near its focus, and the electrons are mostly blown out over the length of the laser. In addition, the ions in (d) are also beginning to drift away from the axis at the back of the laser. At around 1.5 mm in the electron radiograph in (a), a portion of the electron beam is seen to be deflected at a large angle away from the central axis. The same feature is also visible between -1 and 1 mm in (c). At this location in the simulation, the plasma electron density in (b) is rapidly decreasing (onset of blowout), so there is a rapid change in the focusing force on the probe. The focusing force increases with decreasing z near the front of the laser, resulting in probe electrons with larger displacements in x. We clearly see in both (a) and (c) a region at the central axis where all probe electrons are evacuated, this likely corresponding to where electrons are blown out. Interestingly, some electrons have collapsed back to the axis at around 1.2 mm in (b), but this appears to have little effect on the probe beam. The structures in the probe beam image become more complicated towards the back of the laser pulse, where ions begin to move and electrons in the low-density region are oscillating about the axis [see -1 to 0.5 mm in (b)]. The plasma dynamics in this region

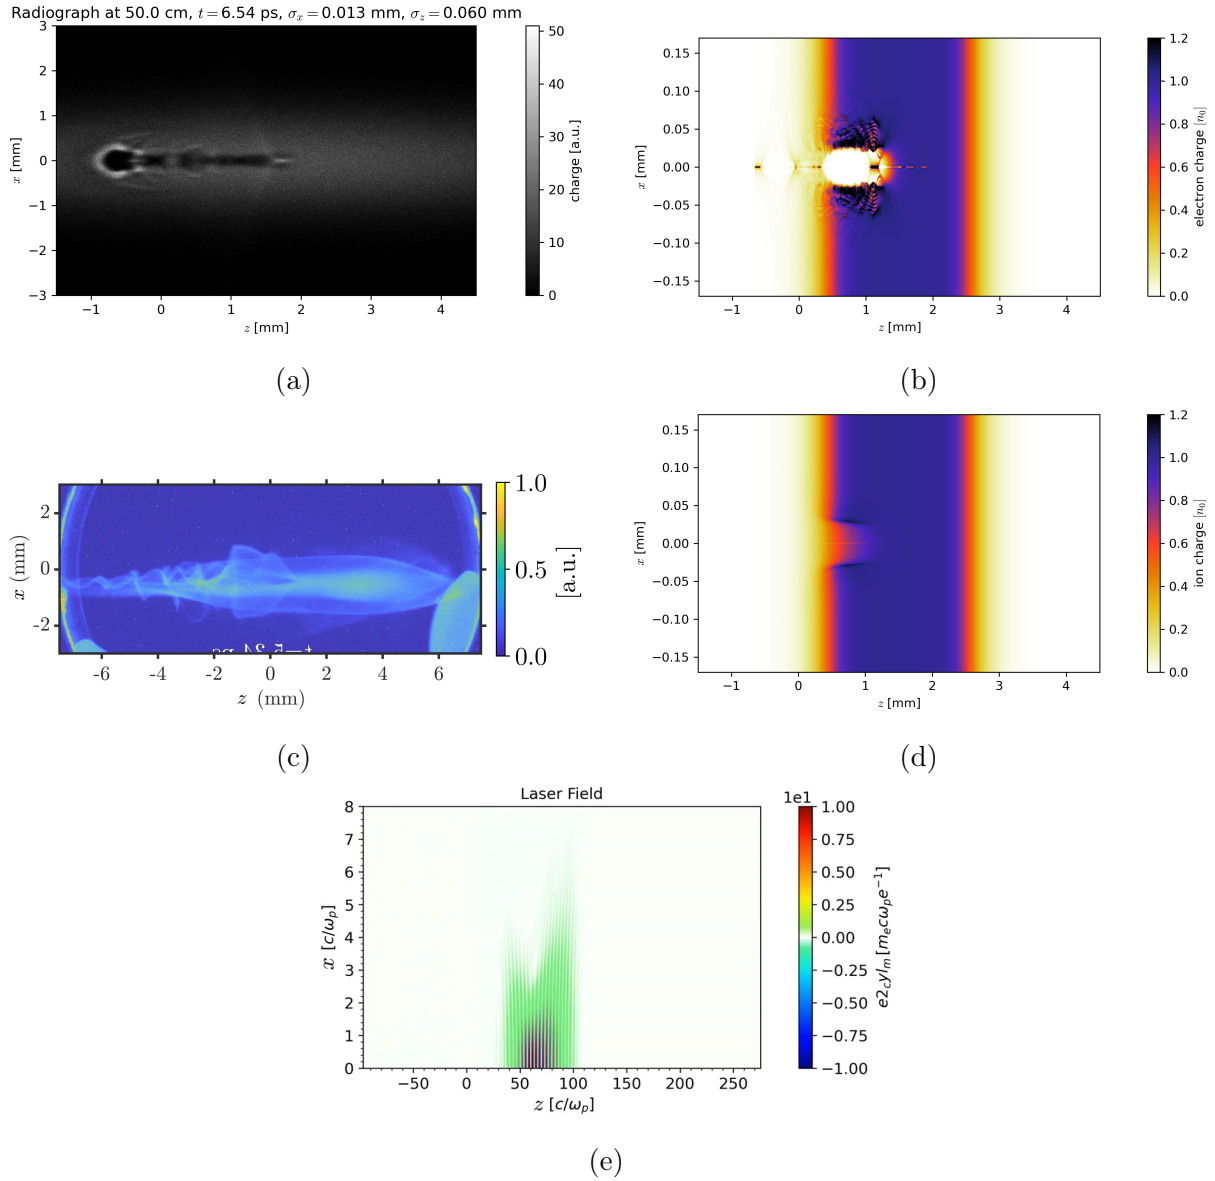

Figure 3.27: (a) Simulated electron radiograph produced at a screen. (b) Electron density exhibiting blowout. (c) Experimental radiograph resembling the simulation results, showing a large "sail" of electrons near the front of the laser followed by oscillations. (d) Ion density, where ions are just beginning to be displaced. (e) Laser field near its focus.

are sensitive to the initial density profile, which may vary from the experiment.

In Fig. 3.28, the laser is past its focus and electrons are no longer completely blown out inside the laser pulse. We see that the structure near the front of the laser pulse is more complicated in both the simulated probe and experiment, although in each we observe a group of electrons whose transverse displacement increases with decreasing z. However, the probe beam features differ significantly between simulation and experiment in the trailing plasma. Note that though there is a clearly defined wake and sideband structure visible in the plasma density between 1–2 mm in (b), this structure is not obvious from the radiograph in (a). This comparison underscores that an electron radiograph does not simply provide a picture of the wake structure, but that more analysis and simulations are necessary to understand their correlation. One issue with comparing the simulated and experimental data is that small changes in the density profile can produce large changes in the radiograph. Although the general shape of the gas density is known from the experiment, this varies from shot to shot.

Figure 3.29 is taken 5 ps after Fig. 3.28, and the laser has almost completely exited the plasma. The (b) electron and (d) ion densities are very similar over the majority of the simulation domain, which means that any focusing force felt by the probe electrons will be much less than before. This is evident by the simulated probe electrons being very close to the axis in the radiograph of (a) between 1 and 2.5 mm (the image plane is likely located near their focus). The experimental radiograph shows similar behavior between -2 and 2 mm. We see at both the right and left edges of (c) some sort of periodic displacement in the probe beam. Similar displacements are seen at each end of (a), but they are much less regular. Once again, the features in the density upramp and downramp are sensitive to the initial density profile.

Finally, Fig. 3.30 is taken 18 ps after Fig. 3.30, and the plasma has drifted away from the axis significantly. Although not directly visible in the plasma density images in (b) and (d), the simulated radiograph in (a) shows two "bubbles" between 2 and 3 mm that appear similar

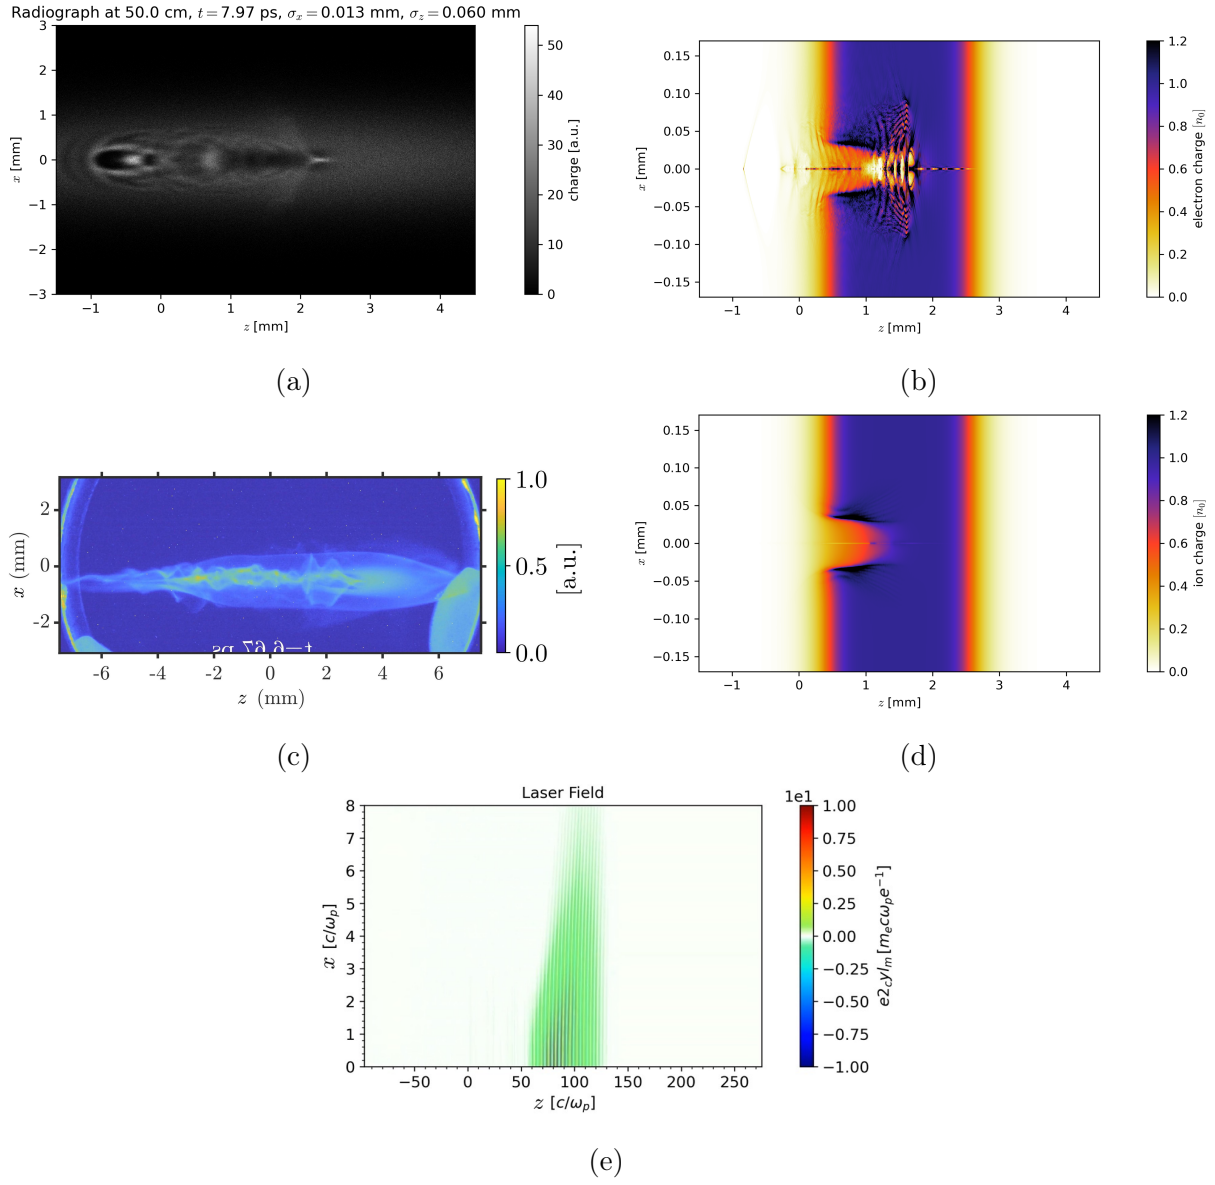

Figure 3.28: (a) Simulated electron radiograph at a screen. (b) Electron density, with several bubbles forming near the front of the laser pulse. (c) Experimental radiograph resembling the simulation results. (d) Ion density, where the electron and ion densities roughly match in the first region of the plasma. (e) Laser fields beginning to defocus.

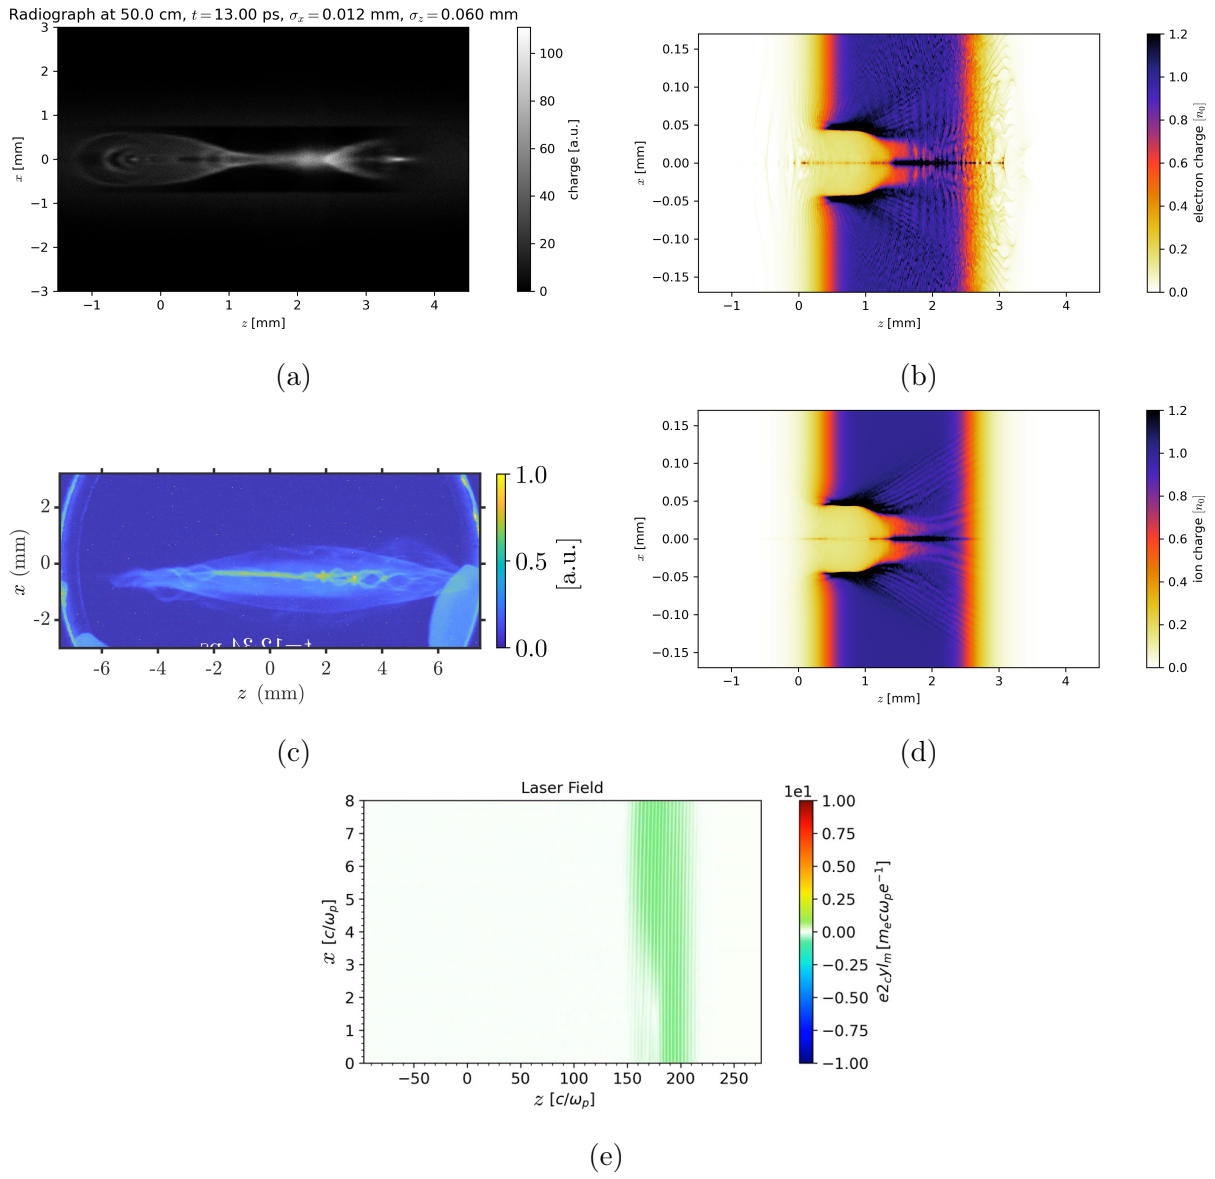

Figure 3.29: (a) Simulated electron radiograph produced at a screen. (b) Electron density with many small perturbations still evolving in the right half of the plasma. (c) Experimental radiograph after the laser has passed through the plasma. (d) Ion density, which matches the electron density over much of the plasma length. (e) Laser fields are exiting the plasma.

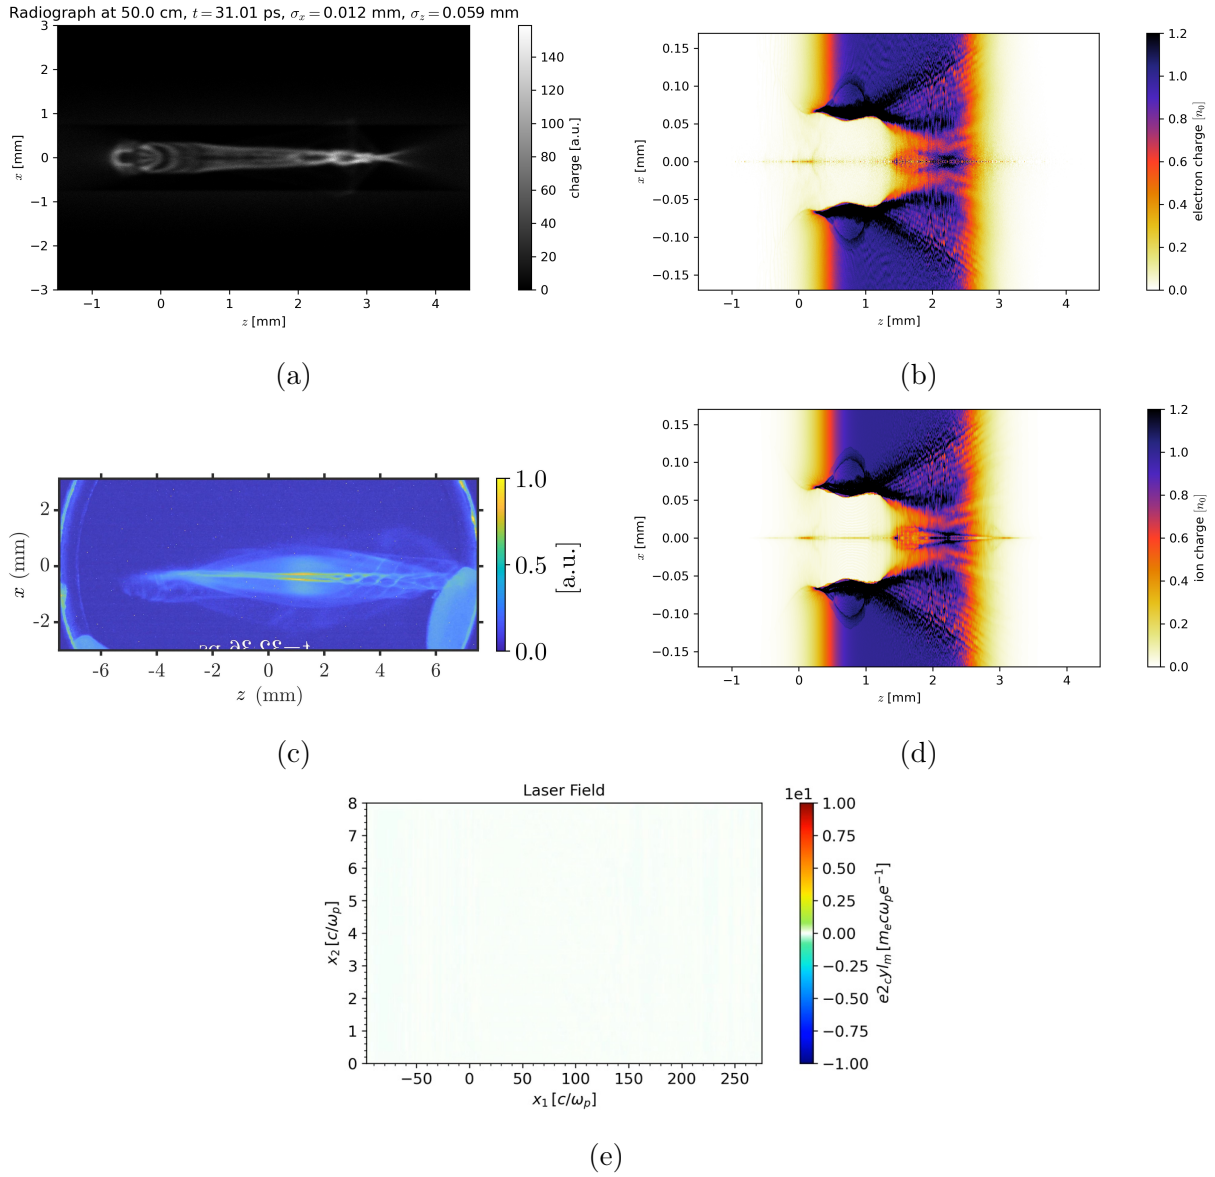

Figure 3.30: (a) Simulated electron radiograph at a screen, with circular features shown at either end of the plasma. (b) Electron density, where the high-density features are observed to be drifting away from the central axis. (c) Experimental radiograph exhibiting circular features at one end and thin features near the axis in the center. (d) Ion density, which almost completely matches the electron density. (e) Laser fields have exited the window.

to those seen in the experiment. In addition, between 0 and 2 mm in (c), we observe two orange lines above and below the axis where electrons are focused. This is also observed in (a), but the data is less accurate because of the finite size of the probe beam and field interaction region (see the outline of background electrons faintly beginning past about 0.8 mm in x).

Some similarities are observed between the simulated probe beam and the experimental radiograph, but we are far from discerning the plasma density from the radiograph. It is likely that in future experiments, the diagnostic screen will have to be moved with respect to the plasma from shot to shot to capture the probe electron motion under variable focusing fields. In addition, a detailed and reliable understanding of the background gas density is crucial to compare against simulation. Another alternative would be to image the wake of a short-pulse laser, where the plasma dynamics are much less complicated. From a simulation perspective, the probe beam could be injected and simulated directly in OSIRIS using the full set of time-dependent fields. Especially late in time, the assumption that all fields drift forward at near the speed of light is problematic and likely results in an incorrectly deformed probe beam. However, overall, electron radiography promises to be an important tool in underdense laser–plasma interactions.

3.5 Summary

Many laser-driven accelerator schemes have been developed that generate energetic beams of electrons. The interplay of direct laser acceleration (DLA) and acceleration from plasma wakefields is an important physical process that is still not yet fully understood. This chapter is devoted to correctly identifying, quantifying and sustaining betatron resonance of energetic electrons in LWFA and SM-LWFA schemes.

We have demonstrated that a picosecond-laser pulse undergoes SM-LWFA in a low-density plasma and that DLA dominates the energy gain of the highest-energy electrons in the absence of a trailing ion channel. This contribution is shown experimentally—and reproduced

with PIC simulations—by the forking structure evident in the dispersed electron beam at high electron energies, as well as through the transition between two temperatures in the measured electron spectra at around 100 MeV. This work provides the first direct experimental characterization, confirmed through quasi-3D PIC simulations with mode separation of fields, of DLA from the SM-LWFA of a picosecond laser. It also provides quantitative comparison between DLA and wakefield acceleration, which gives a better physical understanding for the development of x-ray sources for HEDS experiments.

From the quasi-3D simulation of self-injection in a nonlinear wake excited by a large-amplitude laser, we also demonstrate that energetic electrons can receive a significant contribution of their energy from DLA, even for an LWFA operating in what is referred to as the bubble regime. Accurate calculation of this energy is necessary to understand its importance and to possibly increase its effect. We have shown that neglecting the longitudinal component of the laser field ($E_{z,m=1}$) in the DLA process can overestimate the DLA energy contribution by as much as 50%. This understanding can inform the design of future laser-driven accelerators, where the longitudinal component of the laser might be constructed to be in phase with the electrons and to contribute to energy gain.

A customized finite-difference field solver was designed to combat the errors in the dispersion of light and in the Lorentz force of particles interacting with a laser in vacuum. The real space and phasespace dynamics of accelerated electron beams in an LWFA were shown to vary greatly based on the chosen Maxwell field solver. The traditional Yee solver can incorrectly amplify the transverse betatron oscillations of resonant electrons, yielding incorrect electron energy and angular divergence when compared to experiment. Integration of the equations of motion for representative electrons also gave results consistent with the improved solver. When analyzing single-particle motion exhibiting betatron resonance, we observed that electrons do not readily transition between betatron harmonics, even in the presence of a counter-propagating laser. However, this process remains to be explored in greater detail in future work.

We also investigated the deformation of an electron probe beam propagating across the fields of an SM-LWFA. An experiment was conducted that produced electron radiographs taken over many shots, successively delayed with respect to the laser driver to create snapshots in time as the laser propagated through the plasma. This process was simulated using the PIC method, with plasma and laser fields obtained from fully self-consistent quasi-3D OSIRIS simulations. The probe beam was shown to deform and exhibit many of the general features observed in the experiment, although the varied strength of the focusing forces on the electrons throughout the plasma made direct correlations between the radiograph and plasma density difficult. However, the technique shows promise as a way of measuring and imaging the in situ fields inside the plasma of an wakefield accelerator.

CHAPTER 4

Hot electron and x-ray generation in laser–solid interactions

Particle-in-cell (PIC) simulations have long been used to study the kinetic effects of laser–plasma interactions with overdense plasmas, with applications including the study of novel x-ray light sources [162, 163, 164, 165, 166], generation of mono-energetic ion beams [167], experiments of collisionless shocks [168, 169, 170], transport experiments through warm-dense matter [97] and the fast ignition concept for inertial confinement fusion [93, 171]. The laser–plasma interactions are often simulated for times on the order of picoseconds (1000s of laser periods) and for distances on the order of hundreds of microns (100s of laser wavelengths and 1000s of collisionless skin depths).

The interaction of a high-intensity laser pulse with a solid target often produces a high forward flux of multi-MeV electrons through a variety of physical mechanisms [95, 172, 173, 174, 175, 97]. The "hot" electrons can produce bremsstrahlung x-rays [10, 14], radiation that has been considered for a wide variety of applications [176, 17, 12]. Work has been devoted to optimizing the bremsstrahlung production and controlling its directionality, including the use of preformed plasmas [175, 14, 177], advanced nanowires [178, 179] and cone-shaped targets [180, 181]. In addition, laser absorption has been shown to rise with increased laser intensity and oblique incidence [103, 94], and particle energies greater than predicted by traditional scaling laws have been observed for long-pulse lasers [182].

In this chapter, we describe OSIRIS simulations carried out to effectively determine how to best maximize the dose of 1–5 MeV bremsstrahlung x-rays produced from energetic electrons

generated by high-intensity laser–solid interactions with a given laser energy. We vary the laser pulse amplitude and duration while holding the energy constant, examining trends in electron temperature and x-ray yield. First we discuss the many techniques necessary to accurately simulate such a problem, then show the progression of simulation behavior as the laser is modeled first as a plane wave and then as a finite-width laser in two dimensions. The use of plane waves in two dimensions is referred to as a quasi-1D geometry. We then present data from a large number of large-scale (hundreds of μm), long-time (tens of ps), fully two-dimensional PIC simulations used in conjunction with Monte Carlo simulations to explore the optimal laser intensity for maximizing the 1–5 MeV x-ray dose.

To make such long-time simulations of relativistic laser–solid interactions possible, we employ several novel PIC techniques, some of which are unique to OSIRIS. Electrons (ions) are simulated with 64 (16) particles per cell, but to avoid numerical stopping of energetic particles from the enhanced wakes from macroparticles [183] (see also Sec. 4.1.1), fast electrons are split into two separate particles at certain energy thresholds; we do this at γ values of 1.4, 1.5, 1.6, 1.7, 1.8 and 1.9, such that one electron accelerated to $\gamma \geq 1.9$ is split into 64 smaller particles. To prevent hot-electron refluxing from simulation boundaries, extended particle absorbers [184] (see also Appendix A.3) are employed over the regions $|y| \geq 110\,\mu$m and $x \geq 130\,\mu$m to gradually stop energetic particles without large electric field growth. Low-energy particles with small charge in the absorbing regions are combined with other particles in the same octant of momentum space to reduce computational load. The left vacuum boundary in the x-direction is placed far from the vacuum-plasma interface so as to be causally separated (varied based on simulation time), and dynamic load balancing [43] is performed every 50 time steps. Last, to avoid numerical grid heating and to reduce the enhanced wakes, we use cubic particle shapes.

In this chapter we first provide details on some of the unique algorithmic features in OSIRIS that enable efficient simulations of laser–solid interactions. This is followed by physics results on optimizing the generation of 1–5 MeV x-rays for a laser with fixed energy.

4.1 Techniques for accurate simulation of long-time laser–solid interactions

4.1.1 Macro-particle splitting to reduce enhanced wakefields

In most situations, a particle simulated in a PIC code represents a group of many charged particles with the appropriate charge-to-mass ratio. Each particle has a finite size (on the order of the grid), minimizing two-body particle interactions and thereby reducing collisional effects—even for small numbers of particles per Debye cube (or square). Generally, a simulation exhibits less noise and increased accuracy when more particles per cell are used. Convergence tests can be helpful in determining the appropriate number of particles to use in a given simulation, as long as processes are not growing from thermal noise (in that case the thermal noise also decreases, and the instabilities will take longer to manifest themselves). Due to the large plasma frequency (compared to the laser frequency) when working with an overdense plasma, the smallest spatio-temporal scales are the plasma skin depth and inverse plasma frequency (and not the laser wavelength and inverse laser frequency). Furthermore, as noted above the number of particles per cell must be large to reduce two-body interactions (collisions) so they do not dominate over collisionless absorption processes. This also reduces the related noise levels. Thus the simulation of laser–solid interactions can be computationally challenging. Furthermore, as hot electrons are generated at the vacuum interface and then propagate through the cool background plasma, these macro-particles can exhibit enhanced wakes due to their large charge—compared to individual electrons executing similar motion. The stopping power scales as q^2/m and not as q/m. The macro-particles lose energy to the creation of these enhanced wakefields and slow down as they propagate through the plasma, quantitatively detailed in Ref. [183].

To avoid the enhanced wakes, one could increase the number of particles per cell for the simulation such that the charge of a given macroparticle is sufficiently small. However, this is very computationally expensive. As proposed in Ref. [183], we instead elect to split

macro-particles based on their energy such that their wakes are sufficiently small, but the total number of particles remains relatively low (as determined by the scope of the problem). This is done by creating an additional particle with the same position and momentum as the original, but with half the charge. The charge of the original particle is halved, and then each particle is imparted a small ($\sim 1\%$ of the original momentum) impulse to separate their trajectories. We do this at various energy thresholds, for example, at γ values of 1.4, 1.5, 1.6, 1.7, 1.8 and 1.9. In this way, one electron accelerated to $\gamma \geq 1.9$ is split into 64 smaller particles and will experience an energy loss of $|\Delta\gamma| < 0.25$ from traversing through the entire plasma. We employ this technique for all simulations shown in this chapter to avoid significant energy loss for energetic electrons due to enhanced macroparticle wakes.

4.1.2 Extended particle absorber to reduce hot-electron reflux

When the laser–plasma interaction at the front of a solid target leads to large quantities of forward-going energetic electrons, a large flux of particles will in turn be found leaving the rear simulation boundary. Independent of the particle boundary condition, the exiting stream of energetic particles can be problematic: either sharply removing the current (absorbing boundary condition) or the replacing of particles from a thermal bath (thermal boundary condition) leads to an accumulation of charge, and then to an electric field buildup at the boundary. This strong electric field will generate a return current that is carried by a hot, rarefied electron population (nearly symmetric to the incident electrons) instead of the proper cold, dense population [185]. The hot return current can both modify streaming instabilities that arise in the bulk plasma and modify the laser–plasma interactions at the front surface. To avoid this electron reflux, the plasma may be elongated such that the rear of the plasma is causally separated from the laser–plasma interaction region for some desired duration[186, 187]. In this case a small vacuum region is often placed to the right of the plasma to simplify the particle boundary conditions. However, elongating the plasma introduces extra overhead from simulating (often very particle-dense) excess material.

In an effort to preserve simulation integrity while shortening the simulated plasma region, we use an extended particle boundary condition that sporadically stops particles of certain energies over a defined distance. The first attempt at such a boundary condition was carried out by Tonge et al. [171], where an absorbing region was used in the center of a circular target to model laser–plasma interactions of relevance to fast ignition. In these simulations, a phenomenological energy-dependent damping term was used in the particle pusher. Related approaches were used by May, et al. [97, 188]. In this dissertation, we use a new approach to enable stopping of electrons over an extended spatial region [184]. In the presence of a low-density, hot particle beam or tail moving forward into the plasma, this extended stopping avoids localized charge buildup or current deficiency that occurs when using an absorbing or thermal boundary condition, thus allowing a suitably cool return current to develop in the background plasma over an extended period of time and space. Details on the implementation and characterization of this absorbing boundary condition can be found in Ref. [184] and in Appendix A, and we utilize one or more such absorbers in all simulations shown in this chapter.

4.1.3 Magnetic field growth in quasi-1D simulations

We reiterate that the goal of the laser–solid simulations in this chapter is to determine the electron and x-ray production from the interaction of a finite-size laser with an overdense plasma with a preformed density ramp. Care must be taken when reducing this complicated interaction to a one-dimensional or quasi-1D (2D plane-wave) simulation geometry. For example, some acceleration mechanisms at a vacuum–plasma interface are completely absent in pure one dimension, where the conservation of transverse canonical momentum holds [97]. If the laser in a two-dimensional simulation is polarized in the simulation plane, then the transverse canonical momentum is no longer conserved when the surface ripples or magnetic filaments inside the plasma are formed. Here we examine the generation of long-wavelength transverse magnetic fields that develop and reflect outward-moving electrons in quasi-1D

simulations, necessitating two-dimensional simulations that are large enough to resolve long-wavelength transverse magnetic field modes.

To examine the rippling of the surface and the variability in magnetic field growth in a quasi-1D geometry, we simulate a 1-μm plane-wave laser incident from left to right on an overdense plasma of density $10n_c$ (where n_c is the critical density) with an exponential ramp of scale length $3\,\mu$m. The axial (x_1-direction) simulation domain is $x_1 \in [-2977.6\,c/\omega_0, 942.4\,c/\omega_0]$. The $10n_c$ plasma starts at $x_1 = 0\,c/\omega_0$ and remains constant until the right edge of the box. There is an exponential density ramp of scale length $18.85\,c/\omega_0$ ($3\,\mu$m) that starts at $-173.6\,c/\omega_0$ in x_1, giving a vacuum region of $2804\,c/\omega_0$ in front of the plasma before the plasma expands. The temporal laser profile is a plateau of duration 2.9 ps with smooth transitions to zero of FWHM 0.05 ps on either side, and the normalized laser vector potential is $a_0 = 3.0$. The laser is polarized in the simulation plane and is normally incident on the surface. Simulation cells are $0.2\,c/\omega_0$ square, where ω_0 is the laser frequency ($c/\omega_0 = 159.2$ nm for a $1\,\mu$ m laser), and the time step is $0.141\,\omega_0^{-1}$. The simulation is periodic in the transverse direction, and in Fig. 4.1 we show simulation data where the transverse width is varied from (a) $4\,\mu$m, to (b) $10\,\mu$m and then to (c) $40\,\mu$m. The plots show a zoomed in region in x_1 near the critical density (between -200 and $200\,c/\omega_0$. In each part the top pane shows the transverse electron current (in the x_2-direction), the middle pane shows the out-of-the-plane magnetic field strength, and the bottom pane shows the spatial Fourier transform of the out-of-the-plane magnetic field. All snapshots are taken 0.2 ps after the laser was extinguished.

We observe in Fig. 4.1 that as the transverse size of the simulation is increased, the magnetic field in the low-density plasma region ($x_1 < 0$) acquires a different topology. In particular, in (a) and (b) we observe magnetic field (and current) structures between -200 and $-50\,c/\omega_0$ in x_1 that extend transversely across the entire simulation domain ($k_2 = 0$). These correspond to electron current sheets that perpetually circulate due to the periodic boundary conditions. From the spatial Fourier transform results, we see that the low-k transverse modes are only properly resolved in (c), where the transverse size is sufficiently

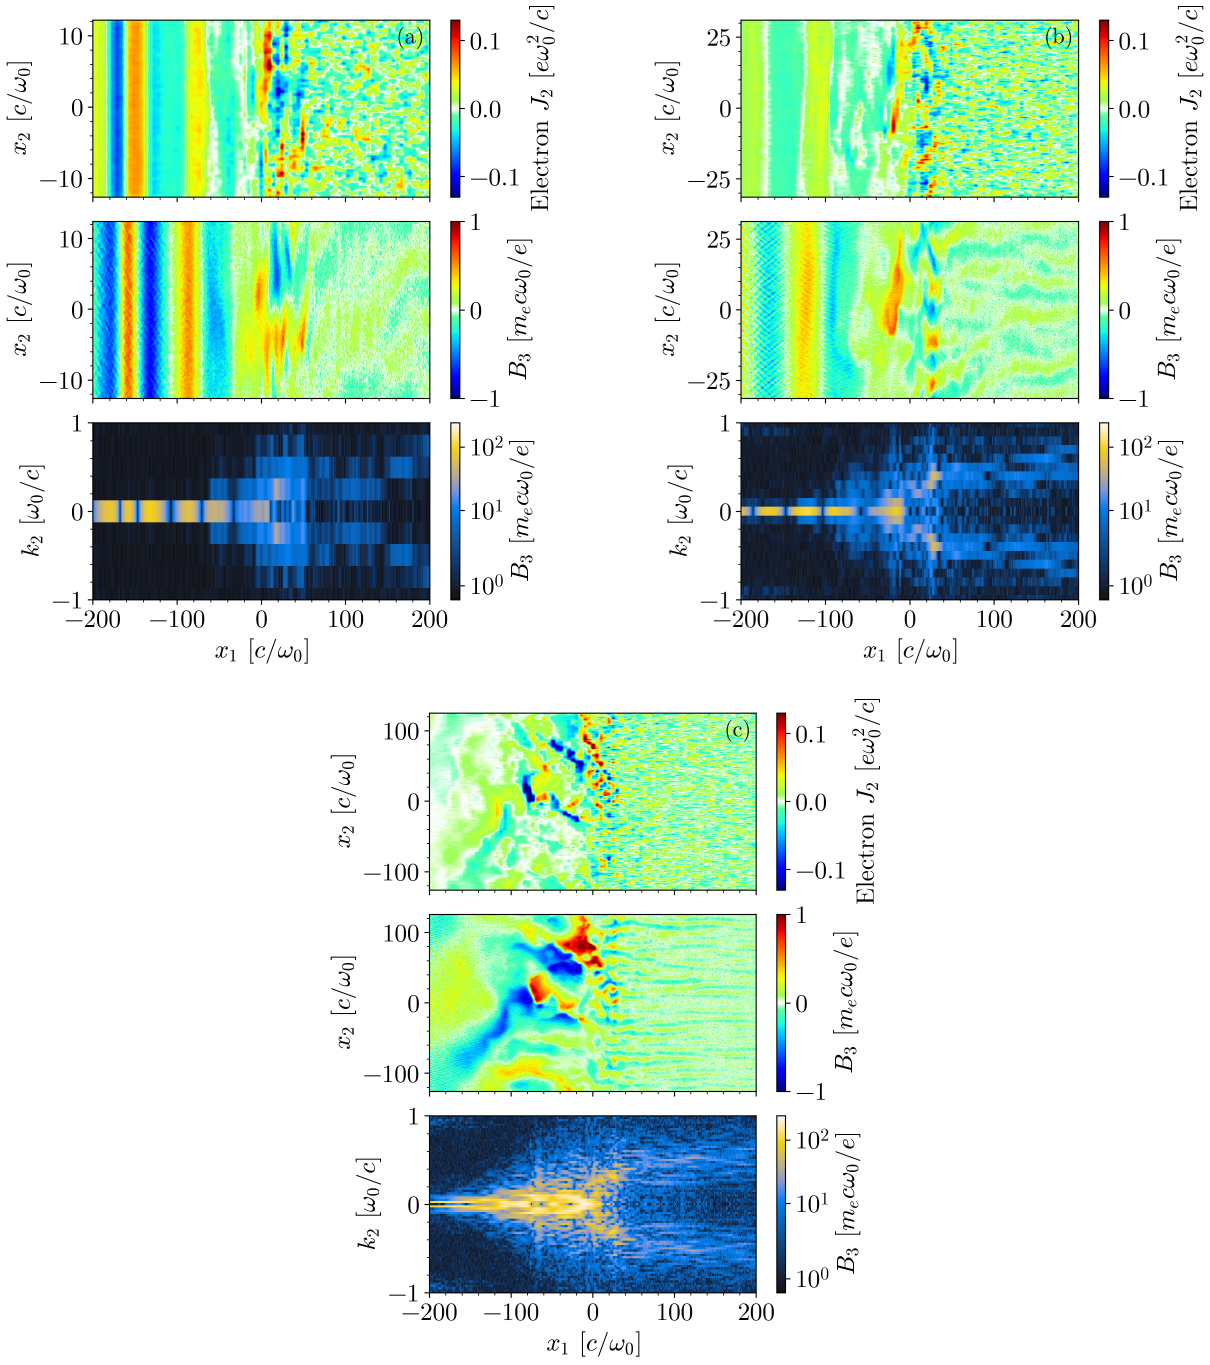

Figure 4.1: Current, magnetic field and spatial Fourier transform of the magnetic field for simulations of transverse width (a) 4 μm, (b) 10 μm and (c) 40 μm (zoomed in for x_1).

large to include at least one wavelength of the largest magnetic field structure. The large periodic currents and magnetic fields seen in these simulations are generated after just 3 ps. Since we desire to simulate a 30-ps laser, it is important to use a simulation box wide enough in the transverse direction to avoid any periodic current sheets that could grow over time.

In Fig. 4.2, we also show the magnetic field and electron phasespace near the end of the simulation (2 ps after the laser was extinguished); these three cases are the same as those shown in Fig. 4.1, but here the entire x_1 domain is shown (full transverse domain shown in both figures). Here the top pane shows the out-of-the-plane magnetic field (the laser field is seen leaving the box to the left), the middle pane shows the p_2-x_1 phasespace and the bottom pane shows the p_1-x_1 phasespace. After the laser is extinguished, the current and field structures convect to the left with the expanding plasma, and they can be seen to completely convect away from the constant-density plasma region. In the middle panes of Fig. 4.2(a) and (b), we can clearly see the alternating transverse current sheets that form the periodic magnetic fields. However, note that these current sheets also correspond to perturbations in the p_1-x_1 phasespace in the bottom panes of (a) and (b), where particles are rotated by the out-of-the-plane magnetic fields. These magnetic fields can act to reflect hot electrons drifting toward the vacuum region from inside the plasma. We also see some acceleration of forward-going electrons in the region $x_1 \in [-1000, -700]\, c/\omega_0$ in the bottom pane of (a).

Perhaps the most important feature of Fig. 4.2 is the amplitude of the p_2 quantities in the expanding plasma ramp for the three simulations. The amplitude of the momentum in each current sheet is larger in (a) than in (b), and in (c) we see that the transverse electron momentum is smoothly spread out over x_1. This is to be expected, since at least one transverse wavelength of the electron current is contained within the box for the larger simulation in (c). If a periodic current/magnetic field structure is left to grow over the entirety of a simulation, the resulting electron momentum can become unphysically large and affect the hot-electron generation over time. The electron temperature between the three simulations shown in Fig. 4.2 are consistent, but the dose of 1–5 MeV forward-going

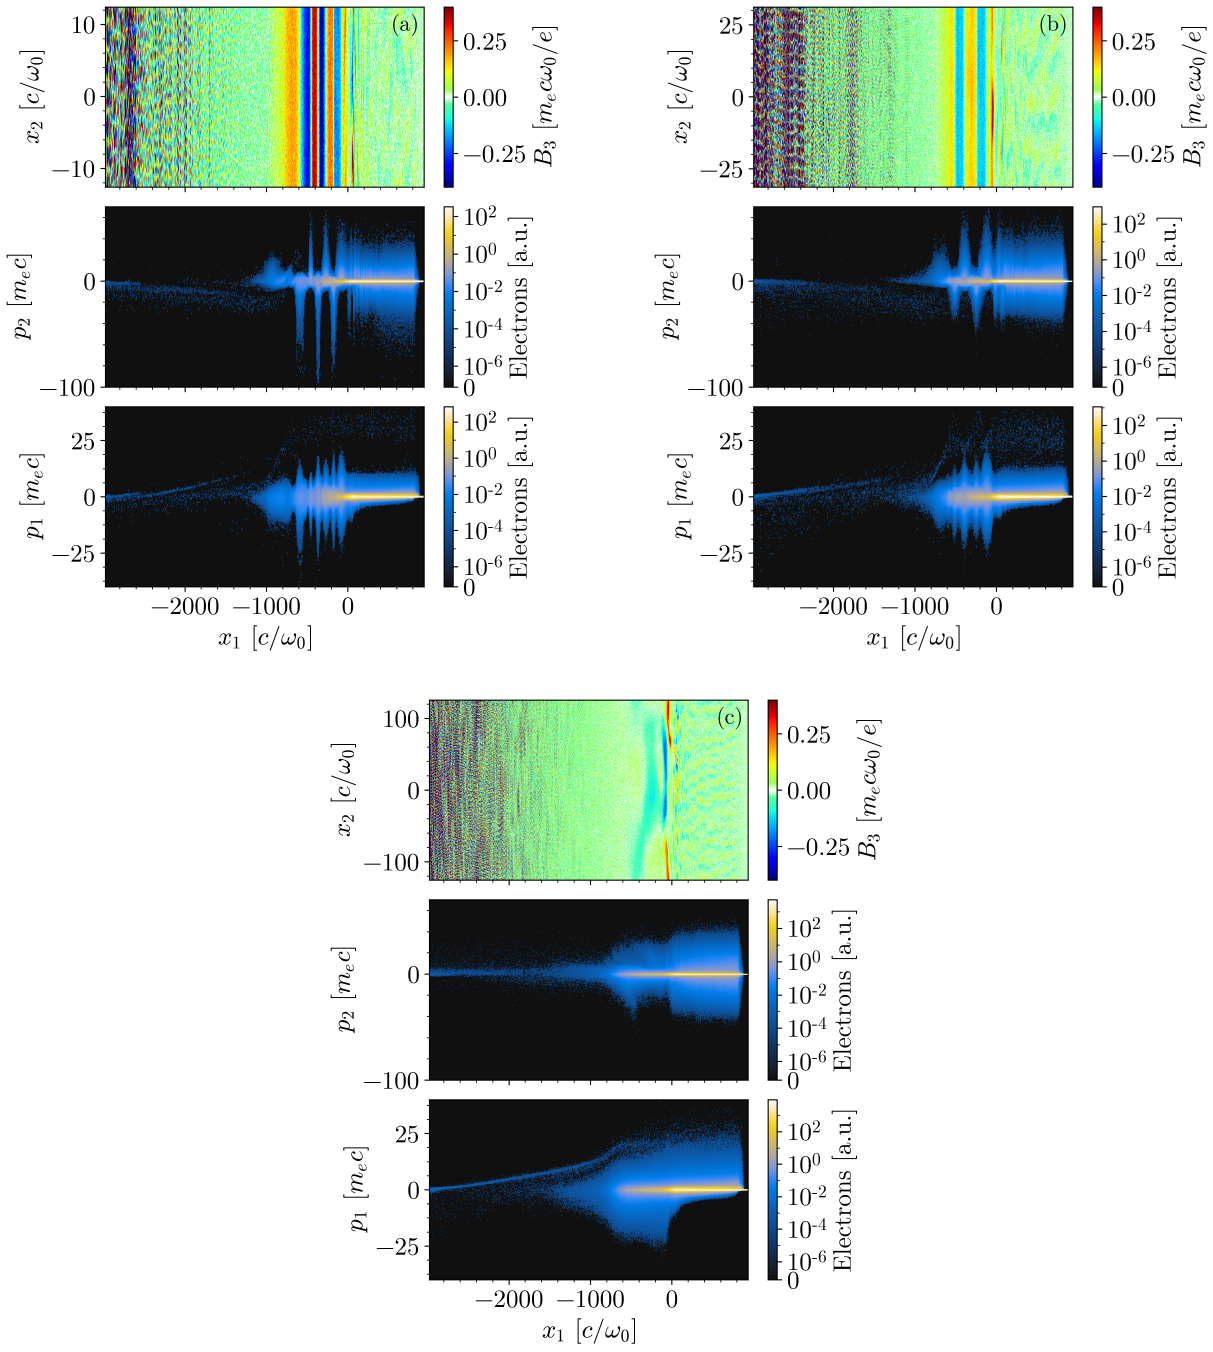

Figure 4.2: View of entire simulation domain showing magnetic field and electron phasespaces for simulations of transverse width (a) 4 μm, (b) 10 μm and (c) 40 μm.

electrons is 6% lower for (c) compared to (a) and (b). Once again, we expect this effect to be further amplified for laser pulses that are longer in duration. It is therefore important to use a simulation box sufficiently wide in the transverse direction to resolve the long-wavelength magnetic field structures generated in the low-density plasma region.

The need for a large transverse simulation size cannot be avoided simply by using a laser polarized out of the simulation plane. In Fig. 4.3, we show in (a), (c) and (e) data from the same simulation pictured in Figs. 4.1(a) and 4.2(a), alongside an identical simulation but with the laser polarized out of the plane, shown in Fig. 4.3(b), (d) and (f). We see in (a) and (b) that fewer energetic electrons are generated for out-of-the-plane polarization; the total number of $> 1\,\mathrm{MeV}$ electrons with out-of-the-plane polarization is only 60% of that with in-the-plane polarization. In addition, the time integrated electron spectra have temperatures of 2.8 and 1.7 MeV for in-the-plane and out-of-the-plane polarization directions, respectively. The surface deformation is also different between the two cases [see (e) and (f)], with the out-of-the-plane case retaining a much flatter surface. There is no real-world analog for the case with a laser polarized out of the plane in a two-dimensional simulation, for canonical momentum is strictly conserved out of the plane. A steep critical surface in such a configuration will not deform or ripple over several picoseconds [188], and many acceleration mechanisms are absent since no heating is possible out of the plane [97]. For this reason we polarize the laser in the plane of the simulation and must retain a sufficiently large transverse simulation size.

4.1.4 Proper transverse boundary conditions for two-dimensional simulations

As described in Sec. 4.1.3, the most accurate simulations of long-time laser–solid interactions require a transverse box size sufficiently large to resolve low-k magnetic field structures that are generated in the low-density plasma. However, for the example explored in Sec. 4.1.3, the largest box size was $40\,\mu$m in width, which is on the order of a reasonable laser spot size for such experiments. Therefore, it is feasible to consider simulating the entire interaction

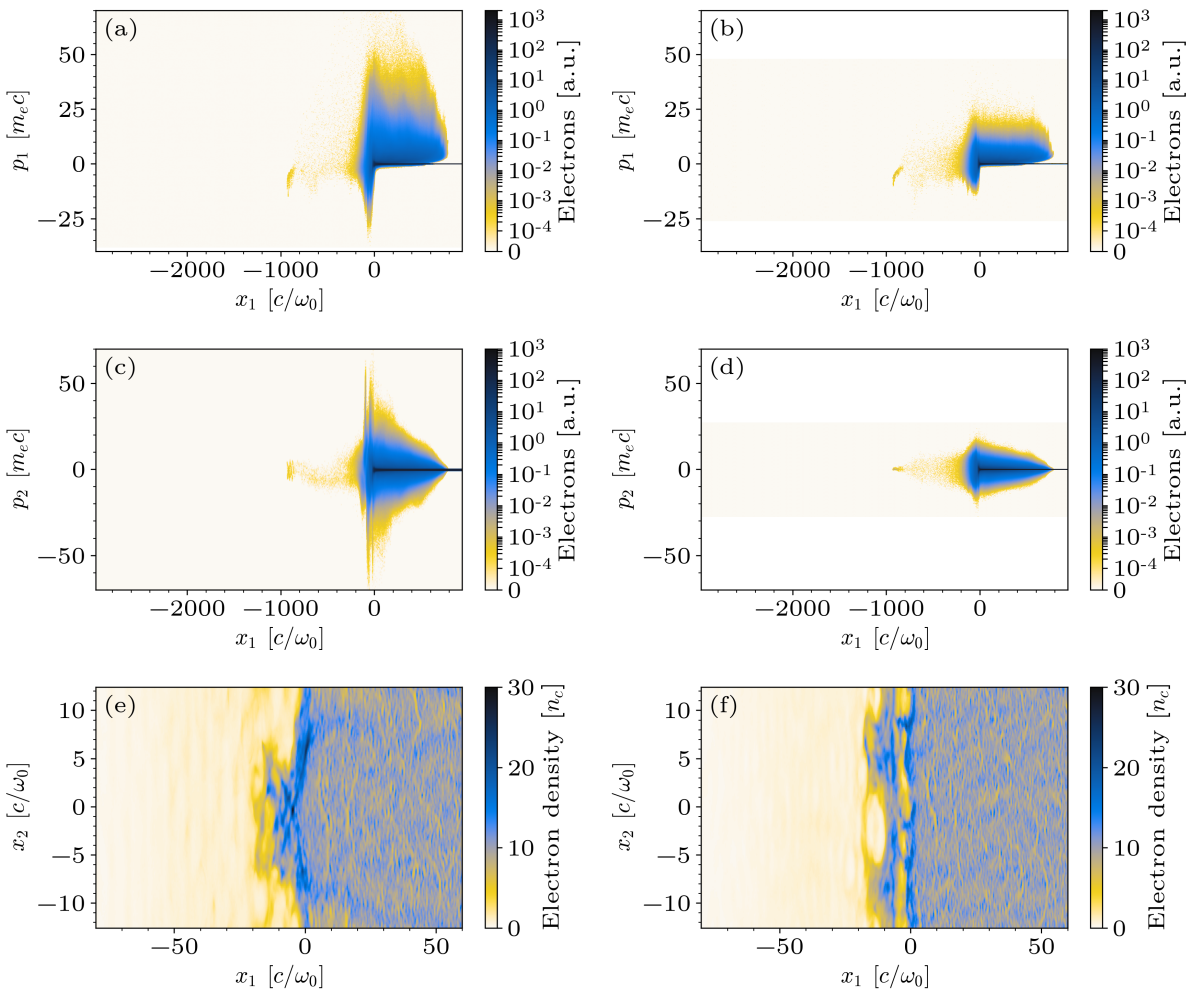

Figure 4.3: Electron phasespace and density plots for the cases with the laser polarized (a), (c), (e) in and (b), (d), (f) out of the simulation plane. When the laser is polarized out of the plane, fewer energetic electrons are generated [compare (a) and (b)], and the surface deformation is very different [compare (e) and (f)].

using a finite-size laser and avoid periodic transverse boundaries. The feasibility of these simulations relies on the assumption that the transverse box size does not have to be large enough to causally separate the boundary from the laser–plasma interaction physics. The challenge is thus to determine boundary conditions in the transverse direction that most effectively reproduce the physics in the interaction region when compared to a causally separated boundary. The boundary condition should limit transverse recycling of electrons when copious amounts of electrons and ions stream away from the center of the laser due to its ponderomotive force. Simply truncating the boundaries can produce the same effects in the transverse direction as described for the longitudinal direction in Sec. 4.1.2 and Appendix A, where a large potential quickly forms at the boundary from the exiting electron currents and produces an unphysically hot return current.

One method for handling the transverse boundary is to put a vacuum gap (after a small density downramp) between the plasma and simulation wall at the top and bottom of the box. This isolated target setup is reasonable if the dimensions of the simulation box are comparable to the target width itself. With this gap, we desire that energetic particles streaming mostly transversely away from the center will drift into the vacuum and either be slowed by the electrostatic Debye sheath formed near the surface [189] or be absorbed at the distant simulation wall. However, the Debye sheath also acts to recirculate electrons back into the interaction region and accelerate ions outward [189, 190]. These processes may or may not be accurate for the desired model setup; if the simulated target is too small, then the use of isolated targets can lead to a large amount of the plasma entering the vacuum region and issues related to refluxing.

Another process associated with isolated targets and finite-width lasers is that there are gradients in density and temperature near the laser spot. This leads to $-\nabla n_e \times \nabla T_e$ magnetic field generation [191, 192], where n_e and T_e are the electron density and temperature, respectively. In addition, electrons expand from the front of the target, creating an electric field normal to the surface. This electric field and the previously described magnetic field

causes electrons to $\vec{E} \times \vec{B}$ drift parallel to the surface [192]. These electrons continue to hug the surface as they reach the top and bottom edges of the bulk density, where they wrap around and begin to propagate in the x_1-direction. As these electrons strike the right simulation boundary, a large potential can develop that accelerates an unphysically hot return current backward in x_1 along the edge of the plasma. This current also wraps around the corner of the plasma density to eventually reach the front of the target, where these electrons can be further accelerated forward into the bulk plasma (producing an artificially energetic spectrum). In this sense, a simulation with transverse vacuum gaps that is truncated in x_1 can behave similarly to a simulation where the plasma extends transversely all the way to the boundary. With transverse vacuum gaps, the energetic transverse currents wrap around the corner of the plasma density and eventually strike (and are reflected from) the x_1 boundary. If the plasma instead extends to the transverse wall, the energetic transverse currents simply strike and reflect off of the x_2 boundary. In both cases, the transverse currents reflux back into the interaction region and generate artificially energetic spectra.

Last, the laser–plasma interactions lead to energetic electrons that stream forward (in x_1) and at an outward angle in the transverse direction. These electrons can exit the plasma into the vacuum regions at the top and bottom of the target. The aforementioned Debye sheath can then reflect them back into the target, while also accelerating ions into the vacuum in the form of target normal sheath acceleration (TNSA) [99, 100]. The presence of the reflected electrons modifies both the spectrum and divergence of energetic particles that propagate through the plasma.

To illustrate some of these effects, we show in Fig. 4.4 the out-of-the-plane magnetic field for a simulation using vacuum gaps between the bulk plasma and walls at the top and bottom of the box (see caption for box dimensions). The plot corresponds to a time after the laser has been extinguished such that no laser field is present. At this time the strong forward flux of electrons is small. To better understand the particle motion and currents that give rise to this magnetic field, we show in Figs. 4.5 and 4.6 the longitudinal and transverse energy flux and

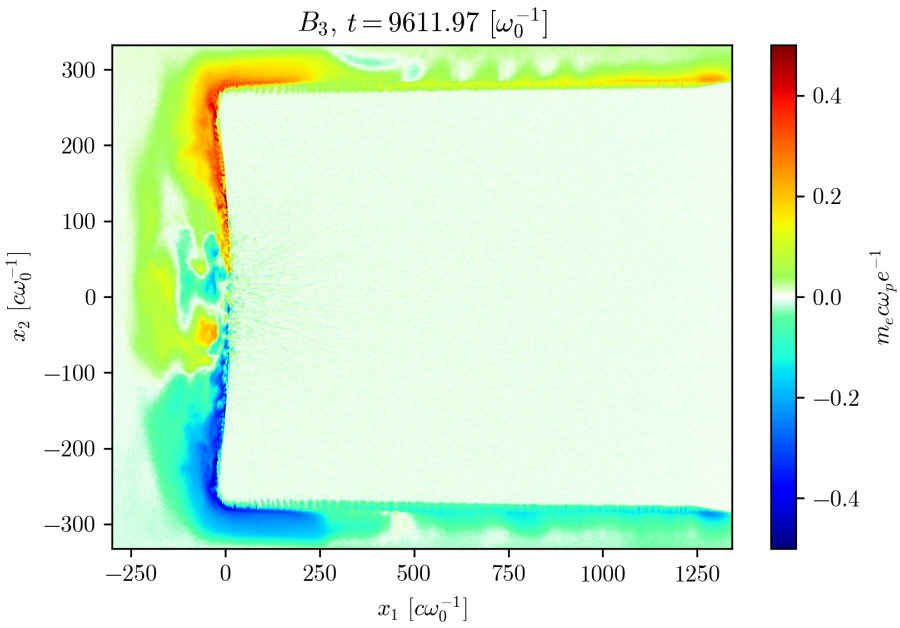

Figure 4.4: The out-of-the-plane magnetic field from the two-dimensional simulation of a finite-size Gaussian laser incident on a solid with vacuum gaps at the top and bottom. The entire transverse box is shown, but the simulation has $x_1 \in [-4329.0, 1342.4]\, c/\omega_0$ longitudinally.

current for electrons and ions, respectively, for a simulation similar to that pictured in Fig. 4.4, but earlier in time (laser is still active) and with larger transverse dimensions (see captions). The energy flux is defined in the i^{th} direction as $Q_i = \int v_i(\gamma - 1)m_s c^2 F_s(\vec{x}, \vec{p})\, d\vec{p}$, where m_s and F_s are the mass and plasma distribution function for a given species, s. The electrons are seen to carry the bulk of the cool return current right at the edge of the plasma—the density decreases from its peak value to zero over the range 124–132 c/ω_0 in the transverse direction. The forward-going electrons [see Fig. 4.5(a) and (c)] are more spread out in x_2 than the returning electrons, resulting in the local magnetic field that quickly decays far from the vacuum–plasma interface (see Fig. 4.4). Note that if these forward-going electrons are not properly slowed before the boundary, a large reflecting potential will develop at the x_1 boundary near the plasma edge and create an unphysically hot return current. The ions carry little surface current, but are seen to mainly move under the influence of the localized laser fields near $x_2 = 0$ [see Fig. 4.6(d)]. We do, however, see a large outgoing energy flux from ions moving into the vacuum region in Fig. 4.6(b) as they are pulled out by the exiting electrons [see Fig. 4.5(a) and (b)].

An alternative to using a vacuum gap in the x_2-direction is to use the extended particle absorbing boundary condition [184] not just at the right boundary in x_1,[1] but also at the top and bottom of the simulation box in x_2. Details are given in Appendix A. In this way, the energetic particles streaming transversely will be slowed while allowing a proper return current to develop. Perhaps the main challenge of this boundary condition is how the absorber behaves in the near-vacuum region in front of the target. The absorber relies on the statistical behavior of many background particles to determine a temperature, but this data is unavailable along much of the x_2 boundary. To overcome this issue, we specify an x_1 value to the left of which a default temperature is supplied in advance; to the right of this value,

[1] If the simulated plasma is also isolated in the x_1 direction (e.g., additional vacuum boundary at the right edge), these surface currents and associated magnetic fields can eventually wrap around the right edge of the target and modify how return current flows backward through the target; this may or may not be accurate for the desired model.

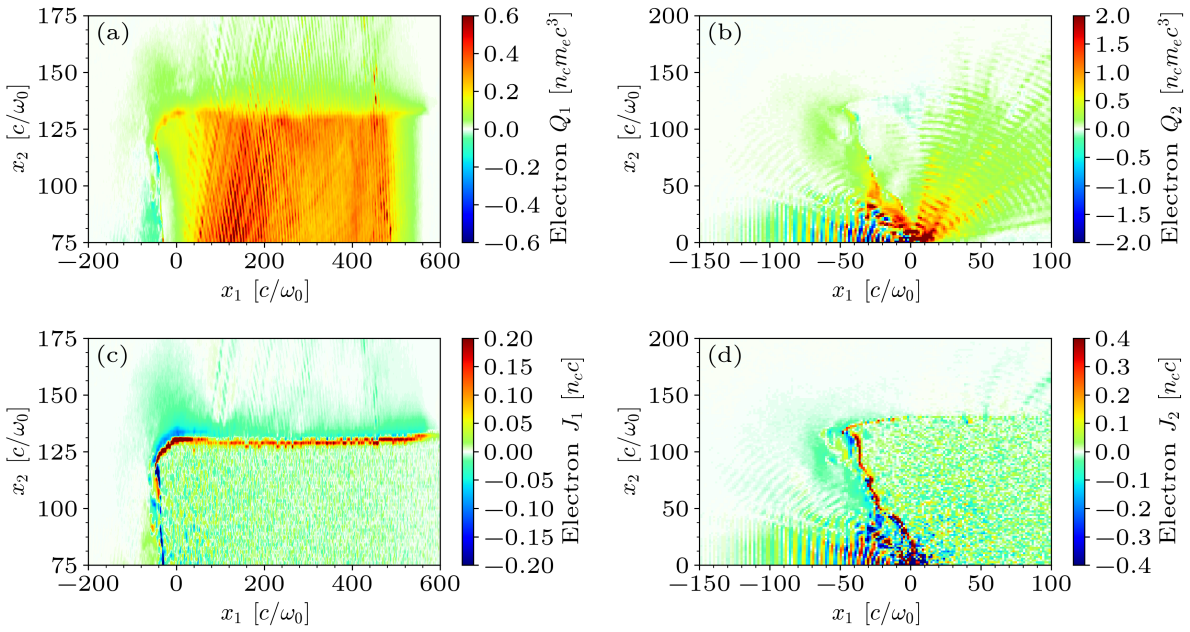

Figure 4.5: Electron energy flux in the (a) longitudinal and (b) transverse directions, as well as the electron current in the (c) longitudinal and (d) transverse directions for a simulation similar to that pictured in Fig. 4.4, but when the laser is still active. The box size is also different, with $x_1 \in [-900, 600]\, c/\omega_0$ and $x_2 \in [-2890.4, 2890.4]\, c/\omega_0$.

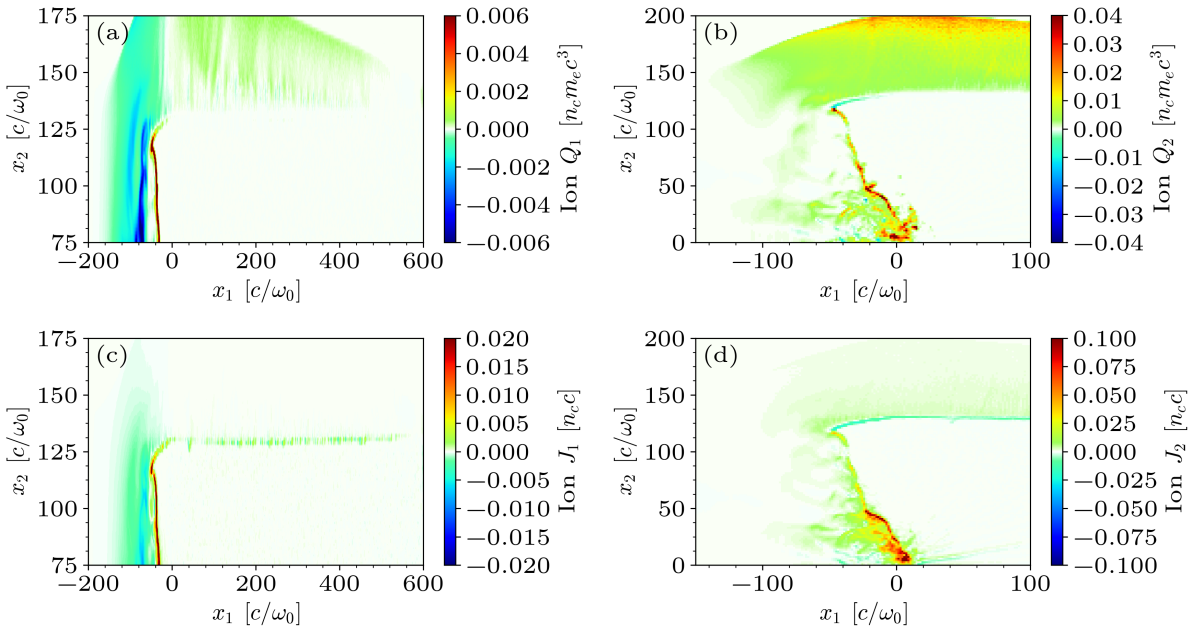

Figure 4.6: Ion energy flux in the (a) longitudinal and (b) transverse directions, as well as the ion current in the (c) longitudinal and (d) transverse directions for a simulation similar to that pictured in Fig. 4.4, but when the laser is still active. The box size is also different, with $x_1 \in [-900, 600]\, c/\omega_0$ and $x_2 \in [-2890.4, 2890.4]\, c/\omega_0$.

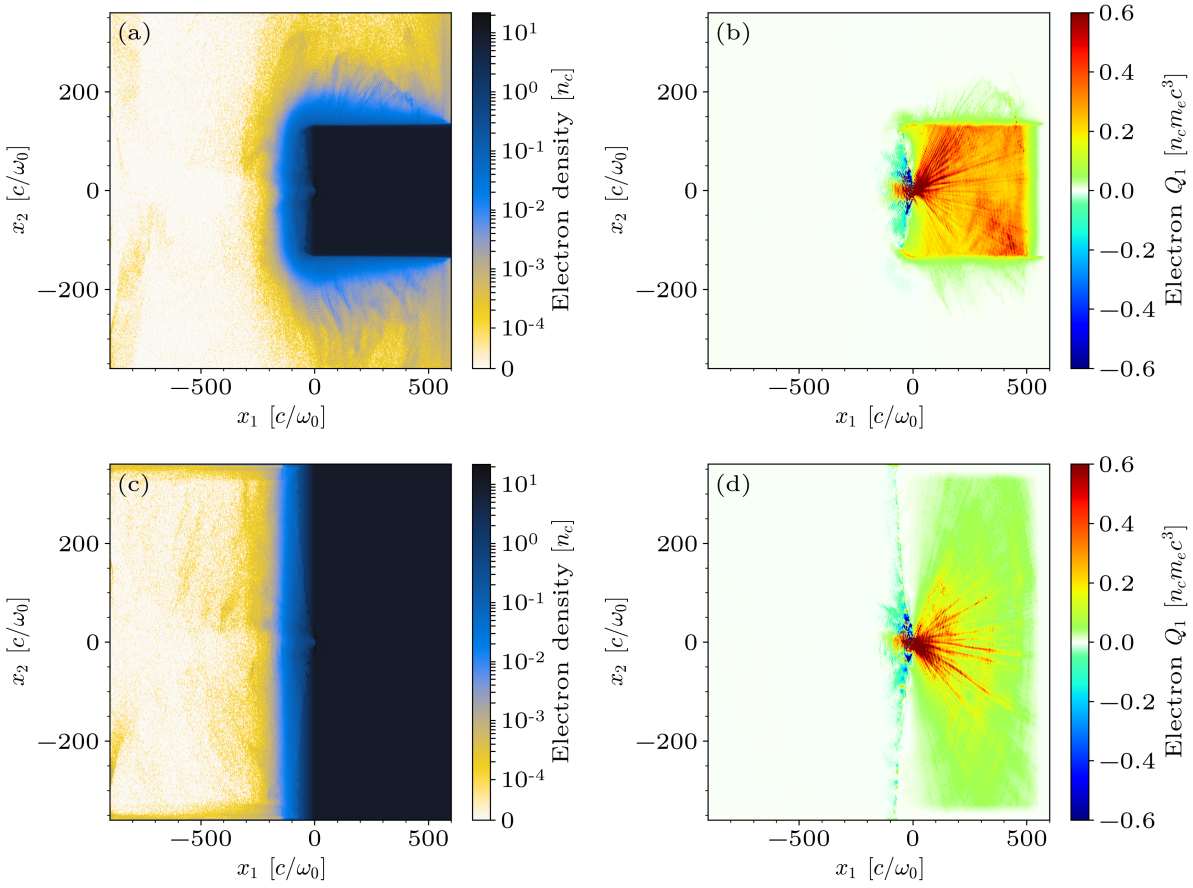

Figure 4.7: Electron density and forward electron energy flux from simulations (a)–(b) containing a vacuum gap in x_2 (transverse box width extends farther than pictured so as to be causally separated) and (c)–(d) using an extended particle absorbing boundary.

the absorber will calculate a temperature from the background plasma. This way particles streaming toward the x_2 boundary in near-vacuum will still be slowed without developing a potential at the boundary cell.

We compare the use of a transverse vacuum boundary (top rows) with the extended absorbing boundary (bottom rows) in Figs. 4.7 and 4.8. All snapshots were taken about 0.7 ps after the laser was incident on the plasma. The entire box size is shown for both simulations except for the transverse width of the vacuum-boundary simulation on the top

row, which extends to $\pm 2890.4\, c/\omega_0$ so as to be causally separated. In Fig. 4.7 we see the electron density in (a) and (c), where transversely going electrons have propagated around the vacuum boundary in (a). We observe the stopping of particles along the x_2 boundary in the near-vacuum region of (c), where electrons begin to collect along the wall. Since particle velocities are slowed in those regions by the absorber, these electrons do not re-enter the bulk plasma. In parts (b) and (d) we see the forward electron energy flux for the two cases. The behavior near the laser–plasma interaction region for (b) and (d) is similar, but in (b) we see that some filaments of energetic electrons are reflected by the vacuum gap in the bottom half of the box near $x_1 = 250\, c/\omega_0$ (as well as at other locations). With the use of the absorber, all hot electrons are slowed within about $50\, c/\omega_0$ of the x_2 boundary, and we only see straight, unbroken filaments of energetic electrons propagating through the bulk plasma.

In Fig. 4.8 we examine the transverse electron and ion energy fluxes for both the vacuum boundary and absorbing boundary cases. We clearly see the reflection of energetic electrons by the vacuum region in (a) from the positive and negative transverse energy flux throughout much of the bulk plasma region. However, in (c) the energy flux is clearly divided along the transverse axis ($x_2 = 0$ line), where the energy flow of electrons is only drifting away from the axis. Finally, we observe from the transverse ion energy flux the presence of TNSA in the transverse plasma–vacuum boundary in (b). Because energetic electrons are exiting into the vacuum, the ions are pulled out and propagate into the near-vacuum region. The current and energy flux in the vacuum regions above and below the bulk plasma thus become important, as they can modify the return flow of energetic particles to the laser–plasma interaction region.

However, from Fig. 4.9 we observe that the transverse return current and dynamics are largely consistent between the cases with and without the vacuum gap. Returning electrons in (a) and (c) are observed to carry the bulk of the return current, whereas the ions in (b) and (d) only exhibit a large current near the laser pulse. Once again, if the extended particle absorber were not implemented at the right simulation boundary in x_1, an unphysically hot

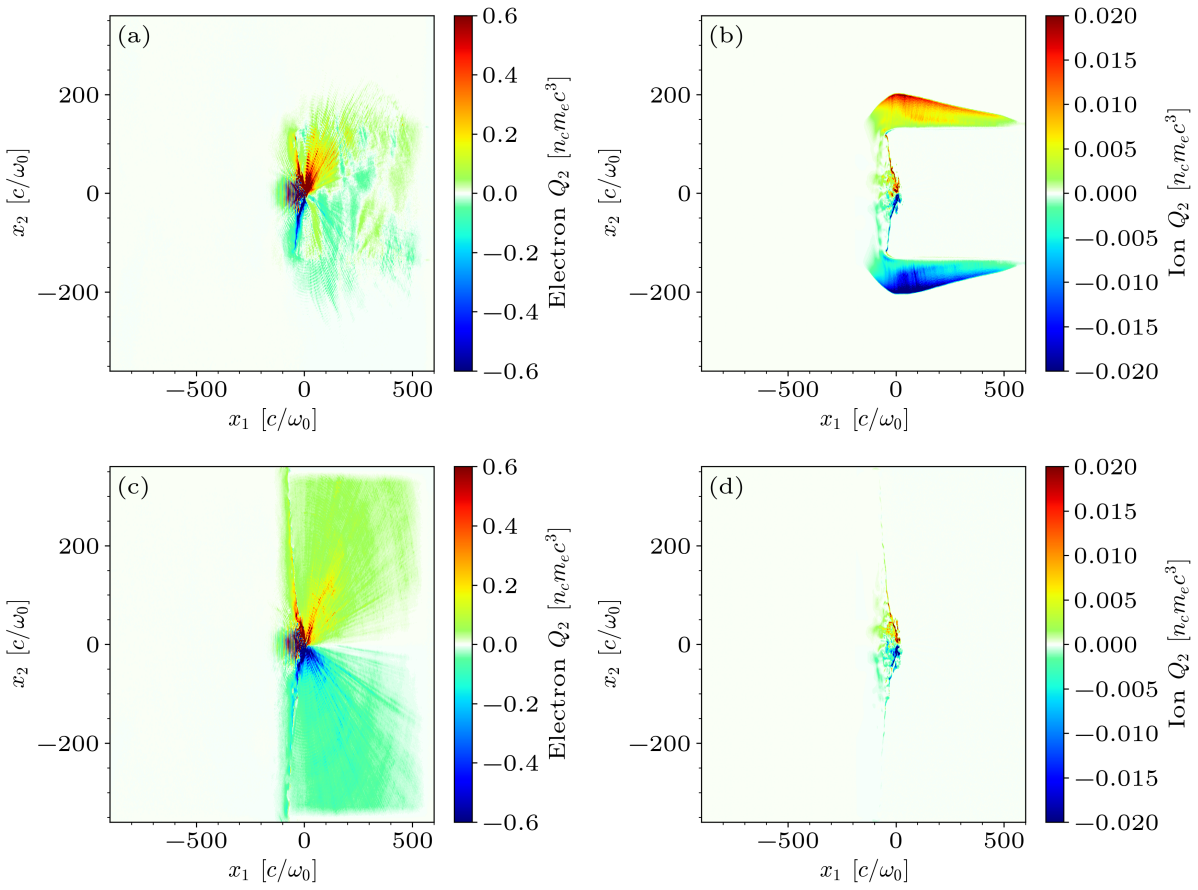

Figure 4.8: Transverse electron and ion energy flux from simulations (a)–(b) containing a vacuum gap in x_2 (transverse box width extends farther than pictured so as to be causally separated) and (c)–(d) using an extended particle absorbing boundary.

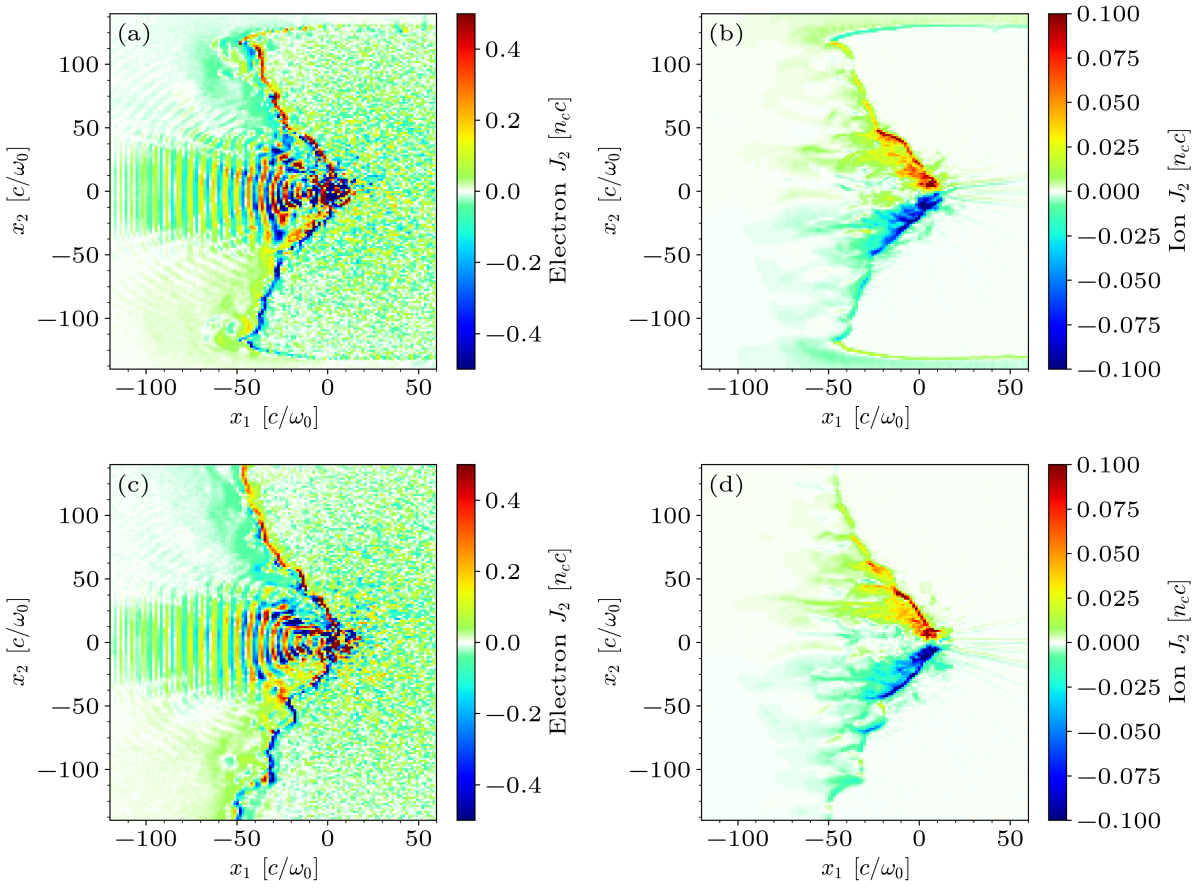

Figure 4.9: Transverse (a),(c) electron and (b),(d) ion current for the simulations (a)–(b) containing a vacuum gap in x_2 and (c)–(d) using an extended particle absorbing boundary. Little difference is observed between the two cases.

return current would flow along the plasma edge and eventually re-enter the laser–plasma interaction region. However, the x_1 boundary in the vacuum gap does not have a bulk plasma to generate a cool return current (even with the absorber), meaning that reflecting potentials can develop and reflect particles backwards in those regions. The electron phasespaces for the two cases are shown in Fig. 4.10, where the hot return current in (a) is located in the vacuum gap (has little effect on the bulk plasma dynamics). The larger values of p_1 for the electrons in the bulk plasma in (a), however, are due to the extra heating that results from energetic electrons being reflected at the vacuum–plasma interface, as seen in Fig. 4.7(b).

From these comparisons, we conclude that the use of an extended particle absorbing boundary in x_2 (as well as at the right boundary in x_1) provides results much more consistent with the behavior of a large plasma that causally separates the transverse boundaries from the laser–plasma interaction region. The transport of hot electrons wrapping around the plasma corner, the reflection of diagonally moving energetic electrons back into the bulk plasma, and the expulsion of ions through TNSA are all problems with using a vacuum gap in x_2 that are avoided by the absorber (the worst of these being the reflection of diagonally moving electrons). The use of the particle absorber makes possible the simulation of fully two-dimensional, long-time laser–solid interactions, which are presented next.

4.2 Parameter scan of laser amplitude and duration

In this section we present particle-in-cell (PIC) simulation results on relativistic laser–solid interactions where the laser pulse amplitude and duration are varied while holding energy constant. As discussed above, several recent computational techniques [183, 184, 43] are employed that allow for large-scale (hundreds of μm), long-time (tens of ps) simulations of such interactions. Electron spectra and beam profiles determined from PIC simulations are used as input for Monte Carlo simulations to calculate x-ray spectra, and a clear maximum in 1–5 MeV x-ray dose is found as a function of laser intensity.

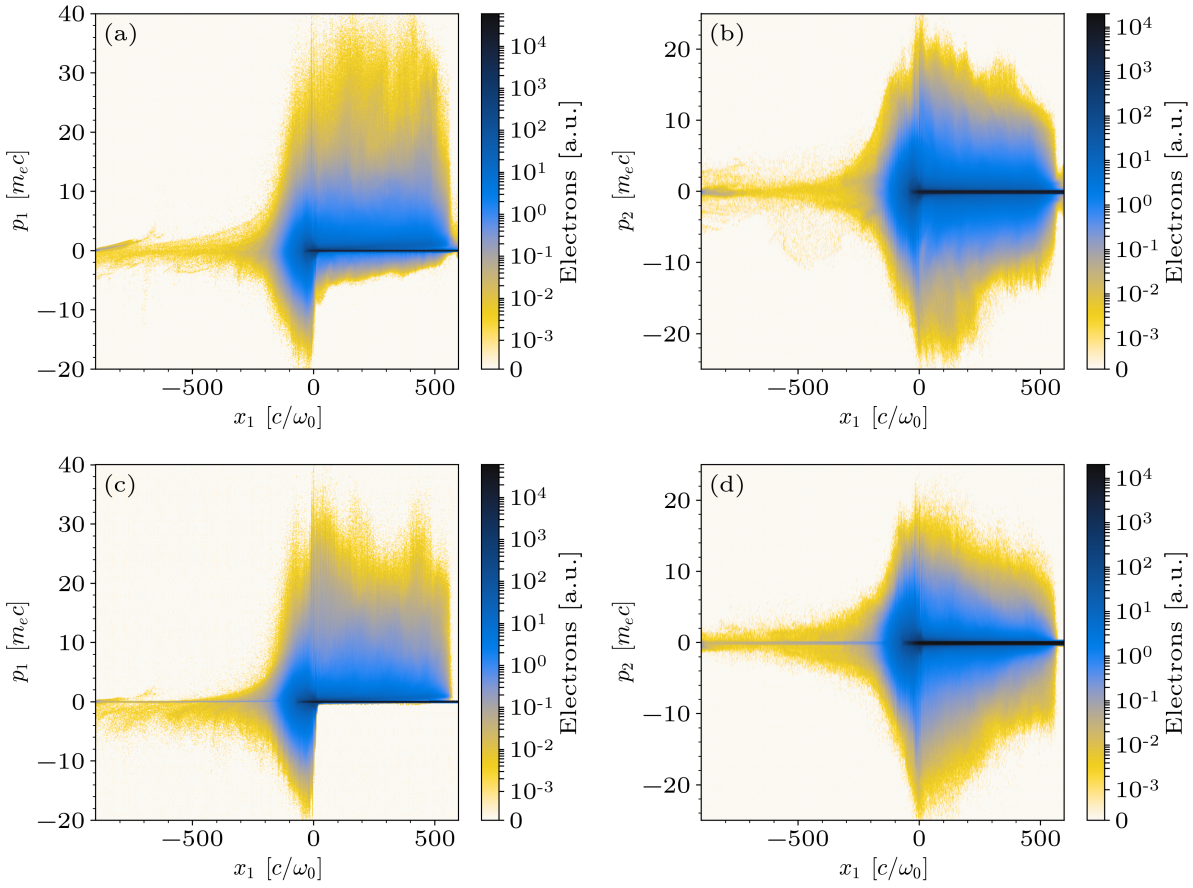

Figure 4.10: Electron phasespace plots for the cases (a)–(b) with and (c)–(d) without a vacuum gap. The hot return current in the bulk plasma in (a) is from electrons in the vacuum gap.

All PIC simulations are performed with OSIRIS [28]. The density rises from near zero to $10\,n_c$ from -27.6 to $0\,\mu$m in an exponential profile of the form $10\,n_c e^{-x/L_0}$, where the scale length, L_0, is $3\,\mu$m and n_c is the critical density; the density remains $10\,n_c$ for positive x. We use two-dimensional simulations where the plasma is $240\,\mu$m wide in y and extends to $150\,\mu$m in x at the right edge. Open and thermal boundary conditions (BCs) are used for fields and particles, respectively, except for an absorbing particle BC at the left wall (and extended particle absorbers near the right, top and bottom boundaries as mentioned above). A diffraction-limited Gaussian laser with spot size, w_0, $30\,\mu$m is launched from left to right with a focus at the critical surface, where the Rayleigh length is $\pi w_0^2/\lambda$ and λ is the laser wavelength. Although each PIC simulation corresponds to a family of cases with identical normalized parameters, we give physical units corresponding to $\lambda = 1\,\mu$m. We vary the amplitude and duration of the laser, keeping the product $a_0^2\tau$ constant at 10 ps, where $a_0 \simeq 8.6 \times 10^{-10}\sqrt{I_0\,[\mathrm{W/cm^2}]}\lambda\,[\mu\mathrm{m}]$ is the normalized vector potential for vacuum intensity I_0 and τ is the full width at half maximum (FWHM) of intensity. The lowest- and highest-intensity cases have $(a_0, \tau) = (0.58, 30\,\mathrm{ps})$ and $(31.6, 0.01\,\mathrm{ps})$, respectively, with energy 200 J for all cases—an energy achievable by many picosecond-class laser systems. The laser profile is Gaussian in the transverse direction and is a polynomial fit to a Gaussian temporally [193]. Cell sizes are $0.4\,c/\omega_0$ in each direction, where ω_0 is the laser frequency. The time step is $0.282\,\omega_0^{-1}$, and simulations end 2 ps after the laser is extinguished. A limited set of simulations are also performed with double the spatial resolution and with a peak density of $30\,n_c$ to ensure that results are consistent for smaller cell sizes and not subject to relativistic transparency. Figure 4.11 shows the electron density, laser envelope and forward electron energy flux (defined as $Q_x = \int v_x(\gamma - 1)m_e c^2 F(\vec{x}, \vec{p})\,d\vec{p}$ for the plasma distribution function F) for a simulation with $(a_0, \tau) = (1.83, 3\,\mathrm{ps})$. The plasma density sharply increases near the laser axis, and the laser is observed to penetrate slightly into the overdense region where the electron energy flux is greatest.

Assuming high-energy electrons deep inside the plasma will continue through the target,

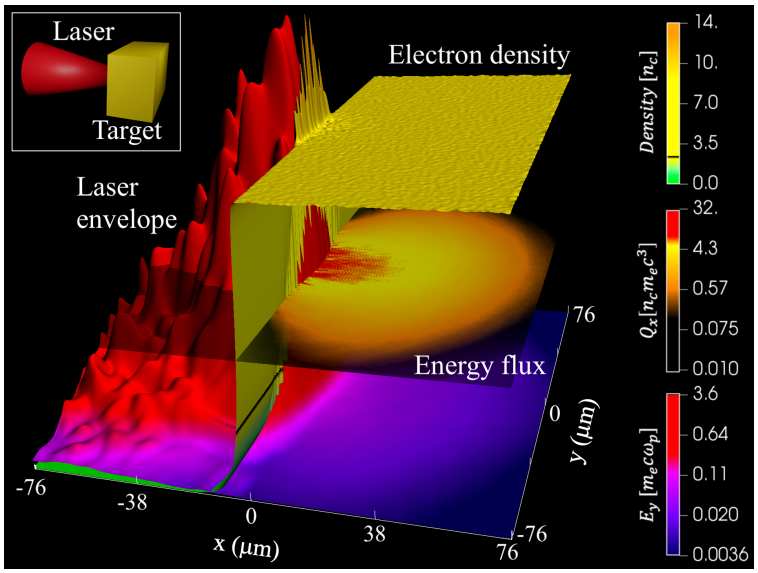

Figure 4.11: Electron density, transverse laser field envelope and forward electron energy flux for a laser with $(a_0, \tau) = (1.83, 3\,\text{ps})$ and w_0 of $30\,\mu$m near when the peak of the laser pulse reaches the critical surface. Quantities are spatially averaged for visualization.

the Monte Carlo simulation package GEANT4 [194] is used to compute the energy and number of bremsstrahlung x-rays generated by these electrons. The GEANT4 code is capable of simulating the bremsstrahlung radiation from individual particles, although it neglects self-consistent electromagnetic fields and other collective effects. The forward-going electron spectra are first extracted from OSIRIS every 200 time steps over the region 50–57.5 μm in x. Next, mono-energetic electrons between the energies of 0.2–100 MeV are injected into 500 μm of tungsten, and the emitted x-ray spectra are gathered as a function of emitted angle. Time-integrated, forward-going PIC electron spectra are then interpolated onto the mono-energetic electron data and weighted to find the corresponding x-ray spectra.

Figure 4.12 shows the cumulative forward-going OSIRIS electron spectra for two simulations. The inset shows the stopping power of tungsten as a function of mono-energetic electron energy. Dash-dotted lines show the resulting x-ray spectra computed from GEANT4 and collected over all angles. Although the low-intensity (blue) simulation in Fig. 4.12 gives

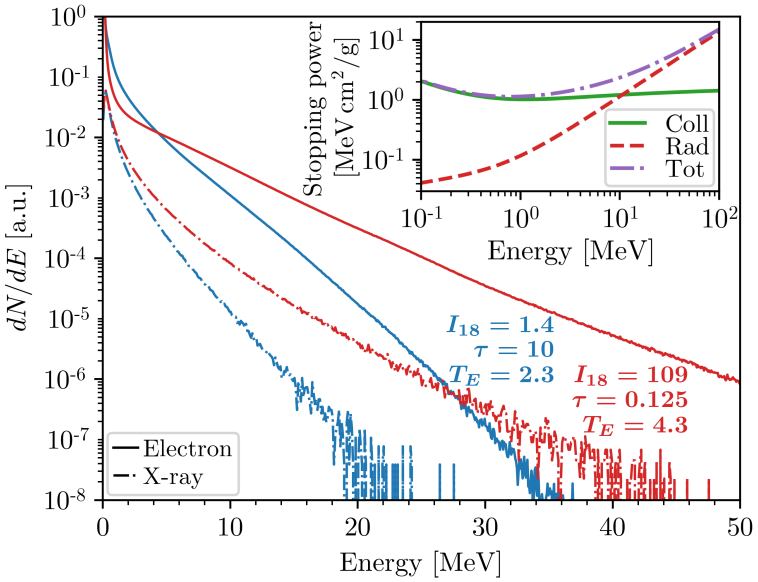

Figure 4.12: Time-integrated electron and resulting photon spectra for two simulations at identical laser pulse energy (laser intensity, FWHM and electron temperature labeled in units of 10^{18} W/cm^2, ps and MeV, respectively). Inset shows the stopping power (collisional, radiative and total) of tungsten as a function of electron energy.

nearly twice the number of 1–5 MeV electrons as the high-intensity (orange) simulation, the large stopping power of many-MeV electrons results in more x-rays of all energies for the high-intensity simulation with a higher electron temperature.

Aggregate data from our sequence of PIC simulations is shown in Fig. 4.13. In (a) we show the amplitude, A, and electron temperature, T_E, of exponential fits done over 3–30 MeV. The generation of energetic electrons is a complicated process that has been described by many differing scaling laws and analyses [105]. Rather than compare to this body of work, as a reference point we also show the ponderomotive temperature scaling [195, 172], which goes as $T_{\mathrm{Pond}} = m_e c^2(\sqrt{1 + a_0^2/2} - 1)$ for a linearly polarized laser. The ponderomotive scaling indicates that T_E scales with the laser field for $a_0 \gg 1$; this corresponds to the work done on an electron after it moves a single laser wavelength, which is more characteristic of the high-intensity laser pulses. However, at low intensity (long duration) the plasma expands, and electrons in the low-density region can repeatedly interact with the laser over a large distance, generating electrons with temperatures and energies much larger than the ponderomotive potential [157, 196, 197, 198, 98].

Insight into how electron spectra are generated and how they scale with intensity can be found by looking at the tracks of individual electrons. The inset to Fig. 4.13(a) shows the distribution of the maximum distance 1500 tracked electrons travel into the low-density region as a function of final energy for the simulation with $(a_0, \tau) = (5.77, 300\,\mathrm{fs})$. On this scale the initial density ramp begins at $27.6\,\lambda$ and becomes constant at $0\,\lambda$, although the constant-density shelf is pushed to $-8\,\lambda$ near the laser axis late in the simulation. For a given energy, there is a distribution of distances that electrons have penetrated into the low-density region, leading to an electron temperature, T_E. However, to obtain an energy of NT_{Pond}, an electron must interact with the laser for at least a distance of $N\lambda$, as seen by the dashed line with a slope of unity. Electrons are colored by the angle of their trajectory with respect to the x-axis. Although no trend in angular distribution is observed with distance traveled into the pre-plasma, we see that more energetic electrons have on average a lower divergence. We

discuss this shortly.

To better predict electron temperature based on the laser and plasma conditions, we look for a temperature model that depends on both the laser duration, τ, and initial scale length of the pre-plasma, L_0. As laser durations and plasma scale lengths decrease to zero, the model should approach the ponderomotive temperature, as observed in a variety of work [95, 199, 171]. We propose a scale temperature

$$T_{\text{Scale}} \propto \left[\sqrt{1 + (\tau/\bar{\tau})^2}\right]^p, \tag{4.1}$$

where the overbar indicates some reference duration and p is an exponent. In this form, the expression equals unity and has zero derivative at $\tau = 0$; we likewise propose the same proportionality to the scale length.

The proposed temperature scaling is thus

$$T_{\text{Scale}} = T_{\text{Pond}} \left[1 + (\tau/\bar{\tau})^2\right]^{p/2} \left[1 + \left(L_0/\bar{L}_0\right)^2\right]^{q/2}. \tag{4.2}$$

Performing a fit of the simulation data to this model yields $(\bar{\tau}, p, \bar{L}_0, q) = (250\,\text{fs}, 0.72, 0.73\,\mu\text{m}, 0.23)$. The fit is plotted as T_{Scale} in Fig. 4.13(a) and is shown to accurately represent the data for all but the highest intensities. The three simulations with highest intensity have lasers with $\tau \leq 25\,\text{fs}$, and many hot electrons are generated when the laser is far from peak intensity. Additional simulations are performed with varied laser energy, E_L, and initial density scale length, L_0, as pictured by the hollow symbols in Fig. 4.13(a). Although T_{Scale} predicts the observed electron temperatures with reasonable accuracy for a wide range of cases, the discrepancy illustrates that the $(\bar{\tau}, p, \bar{L}_0, q)$ coefficients may vary for drastically differing plasma density and laser spot size, wavelength and energy.

In Fig. 4.13(b) we report the time-integrated, forward-going electron dose for three energy ranges (normalized to the largest value shown), along with the fraction of the laser energy absorbed by forward-going electrons. The dose of 1–5 MeV electrons is largest for the low-intensity laser pulses, while the number of $> 5\,\text{MeV}$ electrons steadily increases with intensity.

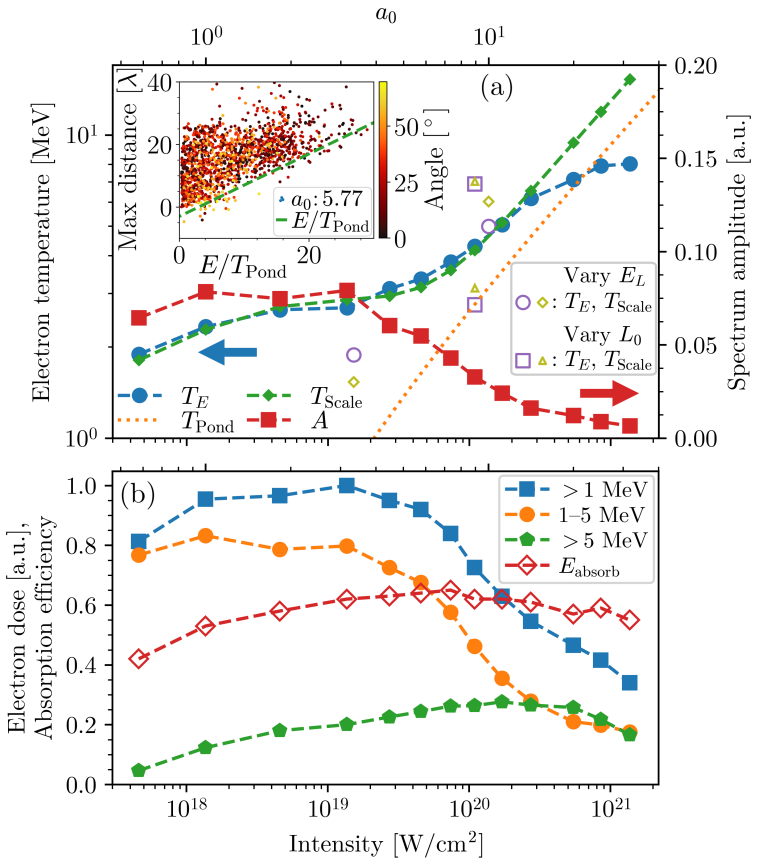

Figure 4.13: (a) Scaling of hot electron temperature with laser intensity at constant energy. The A and T_E values result from an exponential fit Ae^{-E/T_E} to the simulated cumulative forward electron spectra, where E is energy. Inset shows the maximum distance tracked particles extend away from the constant-density region as a function of final energy for a single simulation, colored by absolute angle from the x-axis. (b) Electron dose (normalized to largest value) and fraction of the laser energy absorbed (E_{absorb}).

However, we observe that the total energy absorbed into hot electrons from the laser remains close to 60% for all but the lowest-intensity cases. The near-constant absorption poses the question of how to best maximize x-ray production given the observed electron doses.

To accurately determine the x-ray production using such a large variety of pulse amplitudes, we must first characterize any differences in the angular spread of the electron beam. To accomplish this we track a 10-μm slice of electrons generated from the peak of each laser pulse as it travels through and spreads into the target from 0 to 130 μm. Fitting the forward energy flux near the vacuum interface to a Gaussian profile in the transverse direction ($e^{-y^2/2\sigma_0^2}$), we compute the initial beam width, σ_0. Assuming a beam evolution of $\sigma(x) = \sigma_0\sqrt{1 + (x/\beta^*)^2}$, we then make a fit for β^* using the energy flux as a function of x. Although such a fit is most useful when the incoming energy is carried by a well-defined beam with low energy spread and high energy, it still provides a useful characterization for the beam. We show the parameter fits in Fig. 4.14(a), along with the spot size evaluated at 500 μm. The electron beam becomes significantly less divergent for the high-intensity, short-duration laser pulses, which is favorable for forward x-ray emission.

From the simulated electron spectra and beam profiles, we then compute the observed x-ray spectra using GEANT4 as described previously. In Fig. 4.14(b), we plot the x-ray dose collected both inside a forward cone of 20° aperture as well as over all angles. The energy range of 1–5 MeV is shown alongside all energies, with the latter scaled by 1/3 for visibility. For context, computing the x-ray spectra using only the amplitude and temperature fits shown in Fig. 4.13(a) reproduces the 1–5 MeV x-ray dose to 1% accuracy (5% accuracy for the highest- and lowest-intensity cases). The forward-going x-ray dose very closely follows the > 5 MeV electron dose shown in Fig. 4.13(b), peaking somewhere between 10^{20}–10^{21} W/cm^2. The 1–5 MeV x-ray dose peaks near 5×10^{20} W/cm^2, demonstrating that high-energy (> 5 MeV) and low-divergence electrons are desirable for producing few-MeV x-rays.

To investigate the decreasing electron divergence at high intensity, in Fig. 4.15(a) and (b) we show electron density, time-averaged magnetic field, electric field envelope and forward

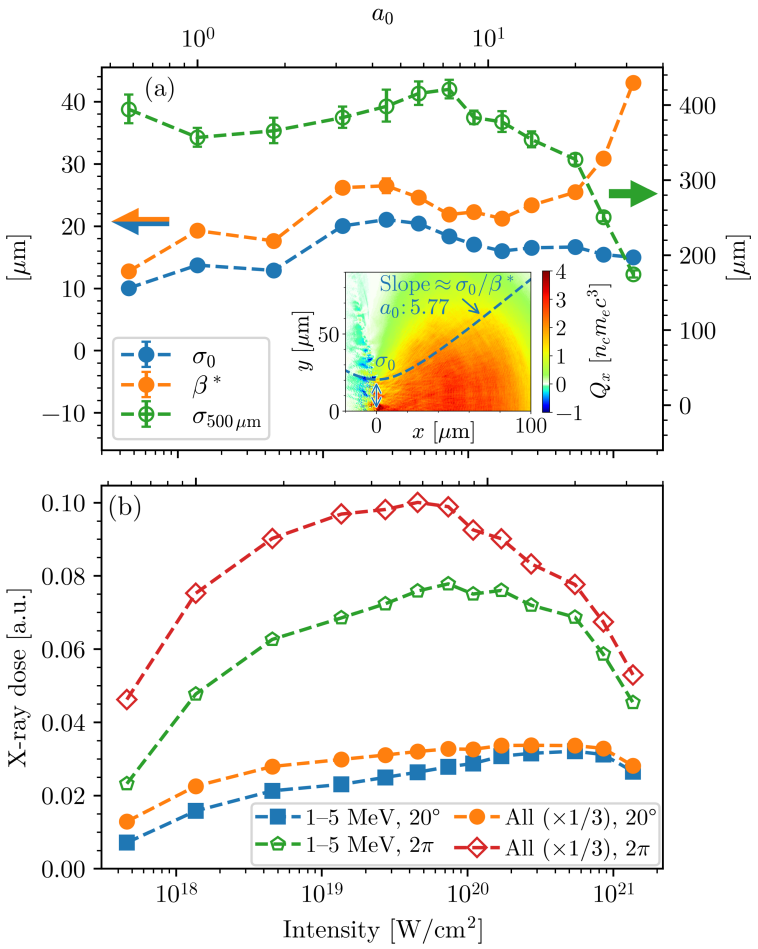

Figure 4.14: (a) Beam parameters and spot size at 500 μm for the hot electrons with respect to laser intensity, where $\sigma(x) = \sigma_0\sqrt{1 + (x/\beta^*)^2}$. Inset shows example fitted beam parameters. (b) X-ray dose as a function of intensity with the same normalization as in Fig. 4.13(b), collected within a forward cone of 20° aperture (blue/orange) and at all angles (green/red). X-ray counts for all energies are scaled by 1/3 for visibility.

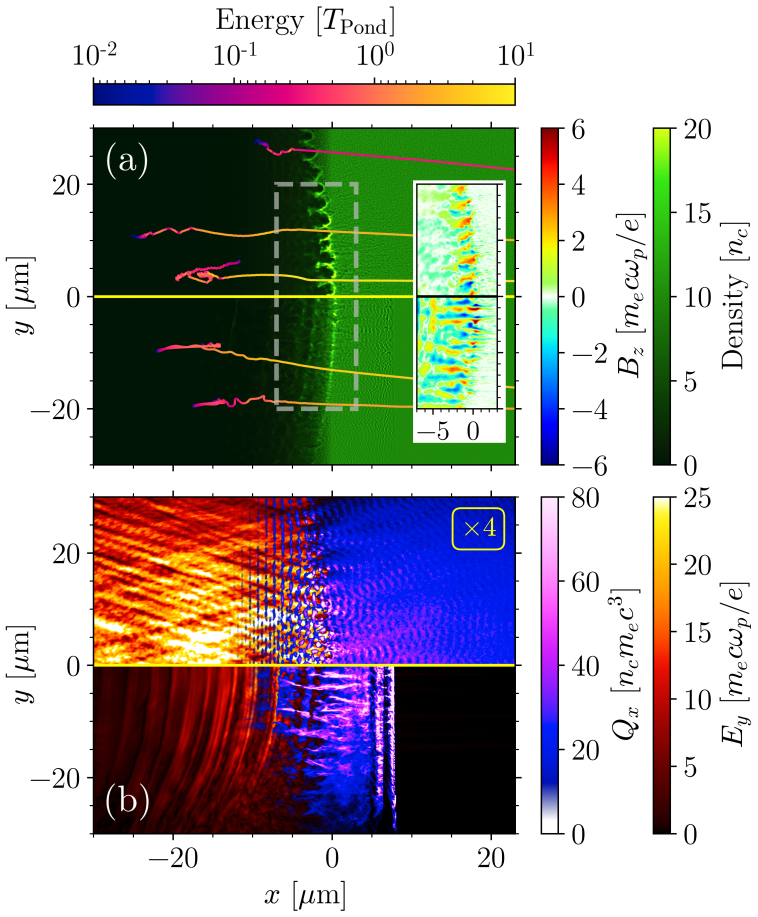

Figure 4.15: (a) Electron density and time-averaged magnetic field (inset for the indicated region) for two simulations, along with the tracks of several electrons. (b) Transverse electric field envelope and forward electron energy flux for the same simulations, with the top pane scaled by four for visualization. In each split-pane figure, the top and bottom halves show the $(a_0, \tau) = (5.77, 300\,\mathrm{fs})$ and $(25, 16\,\mathrm{fs})$ cases, respectively.

electron energy flux for long- (top) and short- (bottom) duration lasers at different times. The far-field electron spot size for the long-duration laser is nearly double that of the short-duration laser. Both panels in (a) show large self-generated surface magnetic fields from a Weibel-like instability [170, 200, 201]. However, in the bottom pane of (b), the first two bunches of electrons (spaced approximately two laser wavelengths apart) are observed to propagate through the plasma with minimal divergence. The maximum cold-fluid-limit growth rate of the Weibel-like instability dominated by oblique modes [202] is given as $\Gamma = \sqrt{3}/2^{4/3}\,(n_b/\gamma_b n_p)^{1/3}\,\omega_p$, where n_b and γ_b are the density and relativistic factors of the beam, and n_p and ω_p correspond to the bulk plasma. The simulation for the bottom panels of Fig. 4.15 has $(n_b/n_p, \gamma_b) = (6 \times 10^{-4}, 12)$, yielding a growth rate of $\Gamma^{-1} \approx 7\,\text{fs}$, or about two laser periods. Tracks of individual particles are shown in (a), where we observe that electrons in the top pane are deflected by surface magnetic fields before free streaming through the remainder of the plasma. Overall, the shortest laser pulses create decreased electron divergence since many energetic electrons are generated prior to the formation of self-generated magnetic fields.

4.3 Summary

In this chapter, we have studied the generation of MeV x-rays from intense laser–solid interactions using a combination of large-scale, two-dimensional OSIRIS (PIC) and GEANT4 (Monte Carlo) simulations. Keeping the laser pulse energy constant at 200 J for a 30-μm spot size, we find that the optimal forward dose of 1–5 MeV x-rays results from a laser intensity near $5 \times 10^{20}\,\text{W/cm}^2$ and duration 25 fs. A new scaling of the energetic electron temperature given in Eq. (4.2) includes a dependence on laser duration and initial density scale length. The formation of Weibel-like magnetic fields near the critical-density surface over only a few laser periods increases the electron beam divergence, leading to decreased forward-going x-ray emission. Simulation results for a wide range of laser pulse parameters show that, for

fixed laser energy, high-intensity (and short-duration) laser pulses generally produce more few-MeV x-rays than do low-intensity pulses due to the formation of electron beams with high temperature and low divergence.

CHAPTER 5

Conclusions

In this dissertation, we investigate the electron spectra when sub-ps, 1-μm laser pulses interact with underdense and overdense plasmas. These electrons can be used to produce x-ray pulses with useful characteristics for probing high energy density plasmas. The research also involved developing improvements to the particle-in-cell (PIC) algorithm that were necessary to accurately study the nonlinear dynamics of electron acceleration via laser–plasma interactions in underdense and overdense plasmas. The results explored and presented here have lead to several publications [113, 74, 37, 203, 184]. In this section we summarize key results and discuss ideas for future work.

Chapter 2 begins with a tile-based dynamic load balancing scheme for PIC codes, along with results from the implementation of this algorithm in OSIRIS. Although tile-based load balancing has been previously implemented in other PIC frameworks [42, 106], we utilize a unique hybrid scheme that gives increased performance and scalability through shared-memory parallelization [113]. Previous implementations of OSIRIS utilize shared-memory parallelization within distributed-memory domains [43], and tile-based schemes use shared-memory cores to process tiles in parallel. We combine these capabilities, allowing for either parallel tile execution with single cores or serial tile execution using all cores at once. This algorithm enables us to use nearly an order of magnitude more shared-memory threads while maintaining good performance [113]. The tile-based structure naturally extends to GPU architectures, and a GPU port of OSIRIS is currently under development based on this structure. In Sec. 2.2, we then discuss three modifications to the OSIRIS quasi-

3D algorithm [30] that yield increased accuracy and stability. We first modify the charge weighting of particles within a few cells of the cylindrical axis; this method correctly initializes a uniform plasma and also eliminates both an on-axis density/current spike and electric field error within an electron beam. We then present a modification of the Maxwell field solver for on-axis quantities to ensure that Gauss's law is satisfied at the cylindrical axis, which is observed to eliminate spurious radial electric field growth for the first-order cylindrical mode in laser wakefield acceleration (LWFA) simulations. Finally, we give an implementation of the Mur absorbing boundary condition [128] in cylindrical coordinates that efficiently absorbs outgoing electromagnetic waves at longitudinal boundaries in quasi-3D. For future work, these improvements could be tested for a broader selection of simulations, e.g., for overdense plasmas. In addition, the on-axis charge weighting scheme introduces a small radial electric field for a thermal plasma with stationary ions, and this should be studied further.

The results given in Chapter 3 pertain to the dynamics of electron acceleration in underdense plasmas. Utilizing the separation of modes in quasi-3D, Sec. 3.1 explores the interplay of direct laser acceleration (DLA) and acceleration from wakefields for the cases of a self-modulated laser wakefield accelerator (SM-LWFA) [74] and an LWFA. Being able to take into account all components of the laser field (i.e., transverse and longitudinal with respect to the propagation direction), we show that the longitudinal component of the laser reduces the total energy gain of electrons in betatron resonance. We also show that electrons gain significant energy from DLA both in an SM-LWFA without the presence of an ion channel, and for a short-pulse LWFA scheme where the laser undergoes photon deceleration [49], causing laser energy to fall back in the first bubble and resonantly accelerate electrons at the bubble rear. This algorithm could be applied to many other regimes of LWFA to study the impact of DLA on electron energy. A customized finite-difference field solver is then given in Sec. 3.2, which provides accurate dispersion for laser light propagating in vacuum and eliminates an error in the Lorentz force on particles from a laser due to the time-staggering of electric and magnetic fields [37]. Full-scale PIC simulations of an LWFA using this improved

solver show that electrons overlapping with the laser in the first bucket have trajectories and energy gains that agree better with experimental data. Further investigation of this motion is performed in Sec. 3.3, where the differential equation of motion for representative electrons is integrated and shown to agree with the PIC results using the customized solver. Higher-order betatron resonances are explored using the same differential equation, where it is observed that for the parameters relevant to the simulation of interest [1], electrons do not readily transition between different orders of resonance under the influence of the laser fields alone. In addition, a counter-propagating laser is shown to eliminate the periodicity of electron trajectories, but otherwise does not provide any significant increase in maximum electron energy. Future work could involve investigating additional mechanisms to transition energetic electrons between different orders of resonance, which could further sustain DLA and give higher electron energies. Then in Sec. 3.4 we compare experiment and simulation for an electron probe beam traveling through an SM-LWFA perpendicular to the laser propagation direction. We see general agreement between the observed features of the probe beam at the screen, although it remains for future work to reverse-engineer the wake structure from the probe-beam images. Another direction for future work is to simulate the probe beam particles directly in the OSIRIS simulation (as test charges), which would eliminate the need to make any assumptions about the motion of the field structures and should greatly improve accuracy.

Finally, in Chapter 4 we aim to determine via PIC simulations the optimal laser configuration (i.e., pulse amplitude and duration for a given energy) to generate the largest dose of few-MeV x-rays from a laser–solid interaction. We begin in Sec. 4.1 by detailing the PIC algorithm advances and necessary simulation parameters to begin to simulate such a dynamic system. High-energy electrons must be split into smaller particles to reduce macro-particle wakefield generation [183], forward-going electrons in the overdense plasma must be slowed near the boundary to avoid unphysical electric field growth [184] (also see Appendix A.3), and plane-wave simulations must be on the order of a finite-width laser spot size in the transverse

direction to properly resolve magnetic field structures. With these considerations in place, in Sec. 4.2 we perform two-dimensional PIC simulations with laser parameters varied over three orders of magnitude of a finite-width Gaussian laser interacting with a solid target. Electron spectra and energies are computed, and a new scaling law is determined to predict the electron temperature as a function of plasma density scale length and laser amplitude and duration. Bremsstrahlung x-ray emission is calculated using a Monte Carlo simulation package and exhibits a maximum dose in the 1–5 MeV range at high laser intensities. These findings can inform the operation, design and construction of current and future laser systems for the production of few-MeV x-rays to be used in high energy density science applications. We find that in order for an electron to gain an energy, E, in excess of the ponderomotive temperature, T_{Pond}, it must make an excursion larger than a laser wavelength, λ, from the critical-density surface. There is a linear relationship between E/T_{Pond} and the excursion distance, L. For a given E/T_{Pond}, there is a distribution of excursion distances, with the minimum value being $L \sim \lambda E/T_{\mathrm{Pond}}$. Directions for future work are to perform similar analyses for different values of laser energy, to carry out full three-dimensional and quasi-3D simulations, as well as to more closely examine the mechanisms of electron acceleration near the critical-density interface.

APPENDIX A

Extended particle absorber for efficient modeling of intense laser-solid interactions

The simulations described in Chapter 4 rely on a particle stopping (absorbing) algorithm to maintain integrity. This algorithm was first conceived of and implemented by Joshua May and Frederico Fiuza, but was then formalized, fully described and implemented into the current version of OSIRIS by the author. A paper detailing this process was recently published [184] and is included here for reference.

A.1 Introduction

Particle-in-cell (PIC) simulations have long been used to study the kinetic effects of laser–plasma interactions with overdense plasmas, with applications including the study of novel x-ray light sources [162, 163, 164, 165, 166], generation of mono-energetic ion beams [167], experiments of collisionless shocks [168, 169, 170], transport experiments through warm-dense matter [97] and the fast ignition concept for inertial confinement fusion [93, 171]. The laser–plasma interactions are often simulated for times on the order of picoseconds (1000s of laser periods) and for distances on the order of hundreds of microns (100s of laser wavelengths and 1000s of collisionless skin depths). The thickness of the simulated overdense plasma region varies depending on the physical setup as well as the acceleration mechanism being explored. For targets thin enough that the entire target can be simulated, vacuum regions are often used on either side of the target, consistent with the experimental setup [204, 205, 206, 199].

For thicker targets that cannot be simulated in their entirety, the plasma may be extended to the simulation boundary [207, 208, 96, 195], where an absorbing or thermal particle boundary condition is used.

When the laser–plasma interaction at the front of the target leads to large quantities of energetic electrons, a large flux of particles will in turn be found leaving the rear simulation boundary. Independent of the particle boundary condition, the exiting stream of energetic particles can be problematic: either sharply removing the current (absorbing boundary condition) or the sudden stopping and accumulation of charge (thermal boundary condition) leads to an electric field buildup at the boundary. This strong electric field will generate a return current that is carried by a hot, rarefied electron population (nearly symmetric to the incident electrons) instead of the proper cold, dense population [185]. The hot return current can both modify streaming instabilities that arise in the bulk plasma and modify the laser–plasma interactions at the front surface. To avoid the electron refluxing, the plasma may be elongated such that the rear of the plasma is causally separated from the laser–plasma interaction region for some desired duration [186, 187]. In this case a small vacuum region is often placed to the right of the plasma to simplify the particle boundary conditions. However, elongating the plasma introduces extra overhead from simulating the (often very particle-dense) excess material.

In an effort to preserve simulation integrity while shortening the simulated plasma region, we propose an extended particle boundary condition that sporadically stops particles of certain energies over a defined distance. In the presence of a low-density, hot particle beam or tail shooting into the plasma, this extended stopping avoids localized charge buildup or current deficiency that occurs when using an absorbing or thermal boundary condition, thus allowing a suitably cool return current to develop in the background plasma over an extended period of time and space. Similar ideas to this extended absorbing boundary condition were implemented in previous versions of OSIRIS as briefly described in Refs. [171, 209], where a drag force was used to gradually slow particles. However, these works did not detail

their approach nor show how performance depended on the drag model. For computational efficiency we instead stop select particles instantaneously over some distance, providing here details for implementation, improvements, potential issues and best practices of such an absorber.

The outline of this appendix is as follows. We first discuss in Sec. A.2 the possible issues with truncating a semi-infinite plasma with standard reflecting, absorbing, or thermal bath particle boundary conditions. In Section A.3 we present the concept and design of the absorber, along with the parameters that can be specified. Finally, in Sec. A.4 we present simulation results from the PIC code OSIRIS, where we test the implementation of the absorber boundary condition on a finite target against a semi-infinite (causally separated) plasma.

A.2 Boundary issues in overdense plasma simulations

The motivation for this work is to efficiently model the interaction of a high-intensity laser with the surface of an overdense plasma. This interaction generates copious amounts of relativistic electrons that propagate forward deep into the target. The forward-going electrons lead to a return current of electrons that then interacts with the laser. When investigating how a laser is absorbed into relativistic electrons, ideally only the region of interest need be simulated; such a region might be an underdense or vacuum region in front of the target or a location some distance into the target itself. However, this presumes that the spectrum of electrons—including its currents and heat flux—are the same as if the entire plasma region was simulated. This may not be the case if the boundary conditions at the edge of the simulation do not properly represent the actual conditions.

The preference of only modeling a small region of a larger plasma is more important for higher target densities, which require smaller time steps and cell sizes (compared to the laser period and wavelength) due to increased plasma frequency and decreased scales of

physical interest. Even in simulations that include a larger transport region for the electrons, it may still not be feasible to simulate the entirety of a physical target—for instance, a millimeter-scale solid-density target. For such simulations to be reliable, it is of course necessary that they not be affected by the choice of the simulation boundary location and the associated boundary conditions. Therefore, it is critical to find a boundary condition that mimics the effects of a quiescent plasma of unbounded depth in the direction opposite from laser incidence. Such a boundary condition is useful even for targets not fully described by an infinite thickness, in that the physics may be understood at least to some approximation as a superposition of unbounded and finite-thickness effects.

Depending on the laser pulse length and/or the target size, it is sometimes possible to simulate an infinitely deep quiescent plasma by expanding the simulated plasma to distances causally separated from the interaction—or perhaps half that far, so signals (moving near the speed of light) cannot reach the boundary and return. Such simulations allow us to determine the "correct" physics, against which we will compare our results. However, in practice these simulations are usually impractical. The required plasma thickness also scales with pulse length, so simulation computer time scales quadratically with the pulse length for one- and two-dimensional periodic cases. For finite-width cases, the transverse size may also need to be extended depending on the pulse length. Therefore, a compact target that reproduces the behavior of a larger or infinite target is desirable.

To make clear the need for an appropriate boundary condition, we mention two spurious effects that can occur in simulations of a truncated target. First, we observe that the laser-matter interaction continuously deposits energy into the plasma. Although the details may be complex, we can assume that the energy will somehow diffuse or dissipate deep into the target; if this is not possible (e.g., inhibited by a boundary), the target will heat artificially. As the absorption of the laser has been postulated to be highly dependent on the target temperature in some scenarios [97], this heating can feed back into the absorption itself and greatly affect the overall simulation.

The second effect of an improper boundary is the modification of the plasma distribution function long before any heating of the bulk electrons through an effect known as electron refluxing [210, 211]. Electrons accelerated by a laser generally make up a super-thermal tail in the distribution function, extending to energies far greater than the background thermal energy [95]. These high-energy electrons have low collisionality (even in solid-density targets [212]) and may thus travel nearly ballistically through the material in the absence of strong fields. At a given transverse plane in the plasma, the bulk plasma exhibits a small backward drift to provide a return current. However, at the plasma boundary the return current formation can be more complicated. For finite-thickness targets with a vacuum region on the far side (both in the lab and in simulations), the first electrons to reach the back of the material exit into a vacuum and continue along a ballistic path. However, an electrostatic field can quickly build at the target rear surface; no such field grows inside the target due to the high conductivity of the background plasma. This decelerating field grows in time, eventually causing energetic electrons to be reflected back into the target. We refer to the re-injection of electrons from this electrostatic field as electron refluxing. After this reflection, the electrons again travel ballistically though the target, where they can reach another vacuum boundary region (sides) and go through another reflection; alternatively they can re-enter the laser–plasma interaction region, in which case their large kinetic energy modifies how they interact with the laser, perhaps significantly.

The reflecting electrostatic field responsible for electron refluxing in finite-thickness targets arises due to the adjacent vacuum region, where escaping electrons leave behind a net charge that resides on the surface. We have found that standard PIC particle boundary conditions responding to high-energy particle beams actually exhibit reflectivity similar to a vacuum boundary. In particular, so-called "reflecting," "absorbing," and "thermal re-emitting" boundaries all lead to refluxing particles early in the beam interaction; in each case a strong electric field builds up at the boundary. A specular reflecting boundary condition—where the sign of the momentum perpendicular to the boundary wall is reversed—clearly leads to

refluxing; however, no electrostatic field is developed at the boundary since the boundary itself reflects the particles.

An absorbing particle boundary, in somewhat simplified terms, simply removes exiting particles (and their corresponding current) from the simulation space. However, electromagnetic PIC codes like OSIRIS advance the fields forward in time via Faraday's and Ampere's laws while depositing current such that the continuity equation is rigorously satisfied. Therefore, when a particle's current disappears, its charge is in fact frozen at its last location. In the case of an absorbing boundary, an exiting electron beam will cause a static charge buildup. The bulk of the plasma near the boundary attempts to shield out the boundary field within a few Debye lengths. However, the field continues to build as more current crosses the boundary and can eventually become large enough to accelerate background electrons backward at relativistic energies, leading to a hot tail of refluxing electrons. Vacuum boundaries behave similarly, where ions are slowly driven off the target via what is known as target normal sheath acceleration [99, 100] (TNSA); again, electrons reach the rear edge in greater numbers than the ions which are able to escape.

A thermal bath particle boundary—where particles are re-emitted into the simulation space with momentum sampled from a specified thermal distribution—may at the outset seem to remedy the issues caused by reflecting and absorbing boundary conditions. We find, however, that a thermal boundary fails to reduce the artificial refluxing. First of all, the correct bath temperature is somewhat ambiguous, and we observe that an incorrect choice leads to clearly incorrect behavior. A bath that is too hot will artificially heat the background electrons in the target; one that is too cold re-emits the particles with too little thermal velocity to diffuse back into the box, manifesting errors similar to those of an absorbing boundary. A drifting Maxwellian moving back into the box (attempting to maintain current neutrality) could be used in place of the stationary Maxwellian generated by the thermal bath. However, determining the proper drift and thermal velocity of this modified Maxwellian is difficult. Furthermore, if a "correct" bath temperature and drift velocity were known from

a causally separated simulation, it would be a function of time. Although there are ways to measure effective temperatures (as discussed in Section A.6), the distribution function is not necessarily well-described by single temperatures; thus it is unclear what one should measure or how to specify the bath temperature.

The extended boundary condition we describe in this appendix is designed with the goal of self-consistently generating a neutralizing, drifting background distribution that imitates an infinite plasma as closely as possible.

A.3 Absorbing boundary region

We desire to model a semi-infinite target with a finite simulation. We must therefore find a way to remove the energetic beam electrons from the simulation while providing a self-consistent return current. One way to do this is to stop the relativistic electrons not at a single point (as in the absorbing boundary condition), but over an extended region. The electrostatic field is then dispersed over a sufficient "volume" such that the electric field driving the reflux is spread out and reduced in amplitude. Thus, a larger "volume" of bulk plasma is driven backwards without reflecting the high-energy electrons. In essence, we create an extended boundary condition for the particles. The plasma can then shield out the field more rapidly such that the potential does not build up over time.

To create this extended boundary condition, we denote some region near the back of the simulation box to be the absorbing region.[1] Within this volume, high-energy particles are selected at random and instantaneously thermalized to a given temperature in a single time step. We do this rather than slowing down the particles gradually—arguably more physically correct—because it allows us to stop the particles throughout the absorber volume without any knowledge of the global beam characteristics; it also obviates the need to track particles across

[1]The absorbing region can also be a finite region in the middle of the plasma. This was done in Ref. [171] for integrated fast ignition simulations. In this case the electrons were slowed gradually by a drag force.

multiple distributed-memory processes as they decelerate (more algorithmically complex). The disadvantage of stopping the particles instantaneously is the emission of Bremsstrahlung radiation [213]. However, this radiation is in most cases both poorly resolved on the simulation grid and can effectively be eliminated by higher-order particle shapes and smoothing.

We define a stopping loop as the process of iterating over all particles in the defined absorbing region, calculating the probability that each particle will be stopped, then stopping a particle (i.e., directly changing its momentum) if a randomly generated number is less than that probability. Stopping loops are performed at a defined time interval, t_a (can be one or multiple simulation time steps). At a given position, the probability of particle stopping and corresponding re-emission are controlled by three parameters: an energy threshold, a mean free path, λ, and a re-emission temperature. The energy threshold provides a way to distinguish between background particles (energy is below the threshold, nothing happens) and hot electrons (energy is above the threshold, stopping may occur). The mean free path, particle velocity and time between stopping loops, t_a, are used to calculate the probability of stopping. This should be done such that the mean distance traveled by an ensemble of particles before stopping is given by the mean free path, but that no energetic particles reach the particle boundary. The re-emission temperature is used to define the Maxwellian distribution from which the particle will be re-emitted in each of the three dimensions. For both the energy threshold and the re-emission temperature, we allow the user to specify either absolute numbers or to calculate the parameters as multiples of the local plasma temperature near a given particle. If required, the local plasma temperature is calculated by integrating the absolute value of the proper velocity in each dimension over all particles in a given cell. See Section A.6 for further details on how this is done. A typical value for the energy threshold is six times the local thermal velocity in any direction, and the re-emission temperature is usually just set to be equal to the local thermal velocity in every direction. Below we discuss two possible methods to accomplish the desired stopping.

A.3.1 Hazard function stopping

Let us assume that all the particles we wish to stop are streaming near the speed of light, c, in the positive x-direction. Rather than stopping these particles at a single boundary that is one cell thick (like a traditional thermal boundary condition), we aim to stop these particles one at a time over the length of an absorbing region many cells thick. Thus particles are stopped over time and space. To facilitate this, we can define a desired normalized distribution that specifies the density of particles in space once they are all stopped. This is in fact a probability density function, call it $f(s)$. For example, we could use the function

$$f(s) = 2\sin^2(\pi s), \tag{A.1}$$

where $s \equiv x/L_a$ for a particle that has traveled distance x into the absorbing region of length L_a. Defined in this way, all particles will have stopped after a distance of L_a. We can also consider a parameterization in time as $s \equiv t/T_a$ for a particle that takes a time T_a to traverse the absorbing region (i.e., $T_a = L_a/v_x$ for a particle velocity v_x in the direction of the absorbing region). The associated cumulative distribution function is then

$$F(s) = \int_0^s f(s')\,ds' = s - \frac{\sin(2\pi s)}{2\pi}, \tag{A.2}$$

and the (continuous) probability that a particle will stop in this interval is given by the hazard function $h(s)$, defined as [214]

$$h(s) \equiv \frac{f(s)}{1 - F(s)} = \frac{4\pi\sin^2(\pi s)}{2\pi(1-s) + \sin(2\pi s)}. \tag{A.3}$$

Since we stop particles instantaneously at discrete intervals rather than gradually slowing them over some distance, we assume that a given particle will have been selected to be stopped after at most $N_a = L_a/t_a v_x = T_a/t_a$ stopping loops (within each stopping loop only a small number of selected particles are stopped). Thus the probability of stopping for each hot particle located at position s for the hazard function absorber is

$$P_{\text{haz}}(s) = \begin{cases} h(s)/N_a = h(s)t_a/T_a & s < 1 \\ 1 & s \geq 1 \end{cases}. \tag{A.4}$$

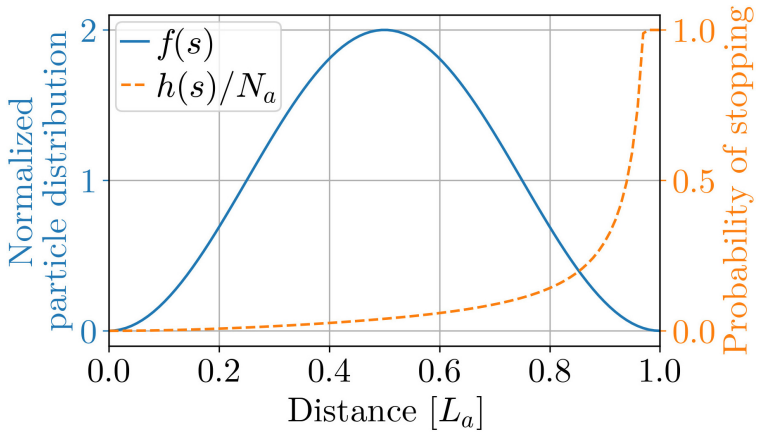

Figure A.1: Normalized probability density function (blue) for desired particle distribution, along with the probability of stopping (orange, dashed) for $N_a = 100$ stopping loops over the absorbing region. Here the mean free path is $\lambda = 0.5\, L_a$.

The probability density function given in Eq. (A.1) is shown along with the probability of stopping for $N_a = 100$ in Fig. A.1. Note that the mean free path for this case is $\lambda = L_a/2$, and thus the scheme can also be parameterized by $s = x/2\lambda$.

Stopping via a hazard function works well when the hot particles are all streaming in the same direction and have large velocity components mainly in one direction (i.e., $|\mathbf{v}| \approx v_x$). However, many simulations employing the absorber may exhibit forward-going particle beams with non-negligible transverse momentum. Since the hazard function absorber stops particles based only on their longitudinal momentum, the background plasma can become disproportionately and unstably hot in the transverse direction. If transversely hot particles have very low longitudinal momentum or are located near the front of the absorbing region, they will have a very low probability of stopping and can continue to drift unchecked for long periods of time (especially dangerous for transversely periodic simulations). In our test simulations the growing transverse currents could eventually form a large reflecting magnetic field near the beginning of the absorbing region, resulting in a reflux of the hot particle beam. To prevent any unstable transverse momentum growth, we propose the following alternative algorithm that stops hot particles without assuming a specific velocity distribution (e.g., all

moving forward near the speed of light).

A.3.2 Linearly varying absorber

A much simpler approach is to stop particles based on their energy alone, independent of their direction of motion or position in the absorbing region. Let us assume that most of the energetic particles we wish to stop are traveling mainly in the forward x-direction at speed c (though some may be travelling backward or have large transverse velocities). Suppose that, similar to Sec. A.3.1, our absorbing region begins a distance L_a from the right simulation edge and that we wish to have a mean free path of $\lambda < L_a/2$ (now in any direction) for the particles, such that they are all stopped before the simulation boundary. An intuitive way to accomplish this is to perform a large number of stopping loops, $N_a \approx 100$, in the time it takes an energetic particle to travel a distance of x_f through the absorbing region; we can naively set the stopping probability to $P = 1/N_a$ for each particle in each loop. We thus require the minimum time between stopping loops to be $t_a = x_f/N_a c$. Such a scheme actually corresponds to a hazard function of $h(x) = 1/x_f$ and a probability density function of $f(x) = \frac{1}{x_f} \exp\big(-\frac{x}{x_f}\big)$. However, after N_a stopping loops (i.e., after particles propagate a maximum distance of $N_a c t_a = x_f$), integrating $\int_0^{x_f} f(x)\,dx$ shows that on average only 63% of particles will have been stopped using the $P = 1/N_a$ probability. We would then need to use the rest of the absorbing region to stop the remaining hot particles. We can extend this idea to construct a more effective absorber that meets our needs.

Based on the above argument, the stopping probability for a particle with velocity magnitude $v = |\mathbf{v}|$ would be $P(v) = v t_a/x_f$. Note that the velocity used is independent of direction, providing equal stopping for particles traveling rapidly forward, backward, or transversely to the absorbing region. However, to facilitate the stopping of particles past a distance of x_f, we parameterize the absorbing region by two longitudinal positions: (1) a location specifying the start of the absorbing region (defined as $x = 0$ here) and (2) a location specifying where the absorber is at its full strength, defined as x_f. Variables such as the

mean free path, energy threshold and re-emission temperature are defined at positions (1) and (2); we will refer to these two kinds of parameters as variables of type "start" and type "full" (x_f refers to the full position). In front of the start position, the absorber is turned off. In-between the start and full positions, the stopping parameters (e.g., energy threshold and mean free path) are changed linearly from their start values to their full values. From the full position to the back of the box, the stopping parameters remain at their full values. The absorber is designed to start with modest stopping parameters and increase to more stringent stopping in order to avoid any sharp transitions in the simulation physics that may result in wave reflections or other spurious effects. Employing a linear ramp comes with the added benefit that we can use a similar set of input parameters for wide range of simulations with varying beam characteristics.

To be specific, we define a scale length, $L(x)$, that varies linearly between the start and full scale lengths, L_s and L_f, respectively, over the distance x_f:

$$L(x) = \begin{cases} L_s + 2(L_f - L_s)x & x < x_f \\ L_f & x \geq x_f \end{cases}. \tag{A.5}$$

The probability of stopping for this linearly varying absorber is then

$$P_{\text{lin}}(x, v) = vt_a/L(x). \tag{A.6}$$

We then let $L_s = x_f$ and define L_f such that particles are rapidly stopped, e.g., $L_f = (L_a - x_f)/20$.

Using this method, instead of initially specifying the probability density function $f(s)$ and then finding the hazard function $h(s)$, we are specifying $h(s)$ first. One can solve for $f(s)$ using Eq. (A.3) and its derivative in s, assuming a piecewise form for $h(s)$. First we write down the hazard function in normalized coordinates as

$$h(s) = \begin{cases} \dfrac{1}{\lambda_s + 2(\lambda_f - \lambda_s)s} & s < s_f \\ \dfrac{1}{\lambda_f} & s \geq s_f \end{cases}, \tag{A.7}$$

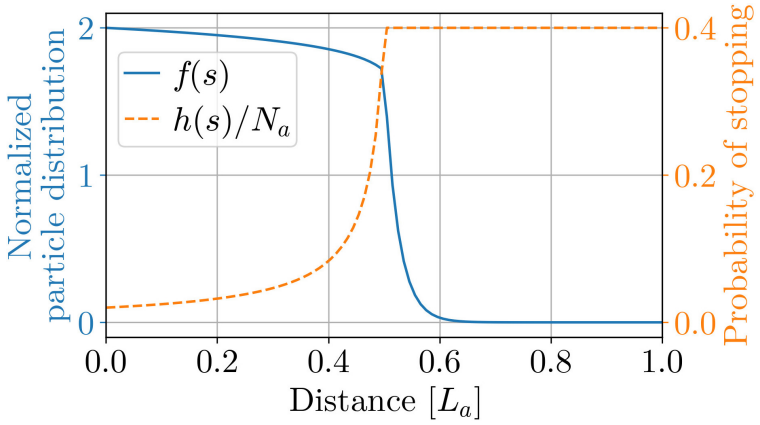

Figure A.2: Normalized probability density function (blue) and probability of stopping (orange, dashed) from Eqs. (A.7) and (A.8) with $\lambda_s = 0.5$, $\lambda_f = 0.025$ and $s_f = 0.5$. Here we have 200 stopping loops over the entire interval, and the mean free path is $\lambda = 0.257\, L_a$.

where similarly to Sec. A.3.1 we have $s \equiv x/L_a$, $\lambda_i = L_i/L_a$ and $s_f = x_f/L_a$. To solve for the probability density function, we require $f(s)$ to be a continuous piecewise function and to integrate to 1, yielding the result

$$f(s) = \begin{cases} s_f \cdot (\lambda_s s_f)^{\frac{s_f}{\lambda_f - \lambda_s}} \cdot \\ \quad \left([\lambda_f - \lambda_s]\, s + \lambda_s s_f\right)^{\frac{s_f}{\lambda_s - \lambda_f} - 1} & s < s_f \\ \frac{1}{\lambda_f} \cdot \exp \frac{s_f - s}{\lambda_f} \cdot \left(\frac{\lambda_f}{\lambda_s}\right)^{\frac{s_f}{\lambda_s - \lambda_f}} & s \geq s_f \end{cases} . \tag{A.8}$$

The probability density function and hazard function for this linearly varying absorber are plotted in Fig. A.2 for $\lambda_s = 0.5$, $\lambda_f = 0.025$ and $s_f = 0.5$. In contrast to the previous scheme, many particles are stopped right at the beginning of the absorbing region (small, non-zero probability of stopping but large number of particles), and the mean free path of the electrons is $\lambda = 0.257\, L_a$ instead of $0.5\, L_a$ for this case. The overall length of the absorbing region and individual parameters can be adjusted to give the desired profile, but care should be taken so that very few electrons ever reach the right boundary. This linearly varying scheme is much more flexible and reliable than the previous scheme. By stopping based on absolute velocity magnitude, particles are stopped whether traveling forward, backward or

transversely to the absorbing region. Specific results will be shown and discussed in Sec. A.4.

A.3.3 Appropriate mean-free-path length

We wish to estimate an appropriate mean free path to use with the absorbing boundary before running the simulation. To do so we will first make an argument based solely on relevant simulation parameters, then make a second based on the strength of the stopping electric field.

First, typical two-dimensional simulations of this type have square cell sizes of length $0.2\,c/\omega_0$, where ω_0 is the laser frequency, with an accompanying time step of $\Delta t \lesssim 0.141\,\omega_0^{-1}$. For simulations such as ours the overdense plasma density is $n = 10n_c$, where n_c is the critical density. If a minimum of 50 stopping loops is desired over a mean free path (100 loops over the absorber length), this sets a minimum value of $\lambda \approx 22\,c/\omega_p$. However, performing the stopping routine at every time step can be computationally expensive; thus stopping every m time steps would increase the minimum mean free path to $\lambda \approx 22m\,c/\omega_p$.

The physics arguments that set the scale of λ can be understood as follows. Collisionless plasmas attempt to remain current neutral. If a net current exists at some instant in time and space, then electric fields are induced on the electron-frequency time scale to neutralize the current. However, this process is more complicated at a boundary. Consider an absorbing boundary; as electrons leave, a net charge builds up, generating a repelling electric field. There is no background charge outside the boundary that can move in to cancel this field. As a result the electric field accelerates background electrons just inside the boundary, growing until it can accelerate electrons backward at near the speed of light. A thermal bath boundary condition results in a very similar situation. Once an electron leaves, another is remitted from a specified distribution function. The probability is chosen based on balancing the flux of a thermal plasma. Therefore, if the electrons leaving come from a hot tail, the remitted electrons cannot properly cancel this current. A large electric field and potential build up, generating a reflux of electrons with relativistic energies. However, if we instead stop the

electrons over a distance along which the plasma can naturally generate a return current, then the physics will most closely resemble reality. The thickness must then be several skin depths wide, i.e., $\lambda \gtrsim \mathcal{O}(10\, c/\omega_p)$, as a skin depth sets the scale over which the current from relativistic electrons is neutralized. This is on the same order of magnitude as what was estimated above for numerical reasons.

Both the numerical and physical arguments predict a mean free path of about the same order, and in Sec. A.4 we explore the performance of absorbers of various size.

A.3.4 Particle splitting and recombination

The absorber is required when simulating overdense laser–plasma interactions because there is a large flux of relativistic electrons moving into the plasma. For such situations there is another issue that must be addressed simultaneously. Due to the relatively large charge (they represent many real electrons) of macroparticles in PIC simulations, these hot particles can generate artificially large wakes, causing them to slow down as they propagate through the plasma [171, 183]. These wakes can also lead to larger levels of turbulence. To avoid the macroparticle-stopping issue, we also run with the particle splitting algorithm discussed in Ref. [183]. This involves defining an energy threshold above which energetic electrons will be split into smaller particles to avoid significant energy loss. Split particles are given a small boost in momentum space (typically 1%) so that they don't travel along identical trajectories.

However, we anticipate that these numerous small particles will eventually be stopped somewhere in the absorbing region. This can lead to a load imbalance in which a disproportionately large number of particles is distributed over a small number of processors, greatly slowing down the simulation as a whole. Since energy conservation is already slightly violated by the stopping of fast particles, this recombination is done simply by averaging the momentum of two particles at a time that lie in the same octant of momentum space. We only look to recombine such particles in the region where particle stopping occurs, and only particles with a charge less than the charge of an unsplit particle are considered for

recombination. Simulations with and without particle recombination give very similar results, but those with recombination run much more quickly (increasing returns for longer-duration simulations).

A.4 OSIRIS simulation results

To demonstrate the effectiveness of the absorber region, we present results from a variety of 2-dimensional simulations of a laser incident on an overdense plasma. Simulations were done using OSIRIS [28], where an absorbing region has been implemented.

A.4.1 Simulation setup

In the simulations, an intense 1-μm plane-wave laser with normalized amplitude $a_0 = 3$ and 3 ps in duration (2.9-ps flat envelope with 0.13-ps rise and fall ramps) is incident on uniform plasma with density $n = 10n_c$ (where n_c is the critical density). The exponential ramp has a scale length of 3 μm and begins at $x = -27.6$ μm. The critical density is then located at $x_c = -6.9$ μm, and we define time $t = 0$ to be when the leading edge of the laser pulse would arrive at x_c if traveling at speed c. The laser is focused to the critical surface and is launched from the left wall. The plasma skin depth is $c/\omega_p = 50.3$ nm and $c/\omega_0 = 159.2$ nm. See Fig. A.3(a) for a schematic.

The simulations used periodic boundary conditions in the second dimension (y), and the laser was polarized with its electric field in the simulation plane (p-polarized). The simulation dimensions were kept constant in the y-direction, 3.2 μm, and in the x-direction were either 923.9 or 1597.8 μm for truncated and causally separated runs, respectively. Square cells of size 0.2 c/ω_0 were used, resulting in a simulation domain of 50197$\times$100 cells for the simulation with the largest length in x (29025$\times$100 cells otherwise). The time step was 0.141 ω_0^{-1}. The electron (ion) species had 64 (16) particles per cell, and each species used cubic interpolation with an initial temperature of 0.1 keV. We employed a static load balancing routine [43] at

initialization to distribute processing elements in an optimal configuration, and the particle push time was delayed until the laser neared the plasma.

In Fig. A.3(b) we show the temporal laser profile, as well as the reflected Poynting flux and transmitted particle energy flux. The reflected Poynting flux is calculated by measuring the total Poynting flux 380 μm before the critical-density interface, then subtracting the known incident laser flux. Both the Poynting and energy fluxes plotted in Fig. A.3(b) are translated in time to line up with the incident laser light. To diagnose the forward momentum and energy flux deep in the plasma, we choose a diagnostic region 48–64 μm into the uniform plasma over which we average the particle data in space. The energy flux is defined as $\int(\gamma-1)m_e c^2\mathbf{p}/\gamma\, d\mathbf{p}$ for electron mass m_e. In order to avoid particle refluxing from either boundary in the x-direction, a 746-μm vacuum region (computationally inexpensive because of the static load balancing) is placed to the left of the plasma upramp, and the uniform-density plasma is extended to the right a distance of 824 μm (computationally expensive). This ensures that any particles reflected from the right boundary region will be causally separated from the diagnostic region for the duration of the simulation (for a time $2\times 760\,\mu\mathrm{m}/c \approx 5$ ps). The p_x-x phasespace is shown in Fig. A.4 at 3.7 ps after the laser was incident on the critical interface, with the diagnostic region marked by dashed lines. Note the large size of the plasma required compared to the diagnostic region location, along with the very hot return current reflecting off the right simulation boundary—even though a thermal particle boundary is being used.

We ran the simulations until 2 ps after the laser had finished hitting the plasma. In all cases, hot particles were split into two after reaching a γ of 1.4, 1.5, 1.6, 1.7, and 1.8 (i.e., very energetic particles were eventually split into 32 smaller particles); the splitting routine was executed every 10 time steps.

The particle acceleration mechanisms in these types of simulations are stochastic; therefore, we expect and indeed do observe large differences in particle statistics due to slightly different simulation configurations. For example, we performed the causally separated simulation three

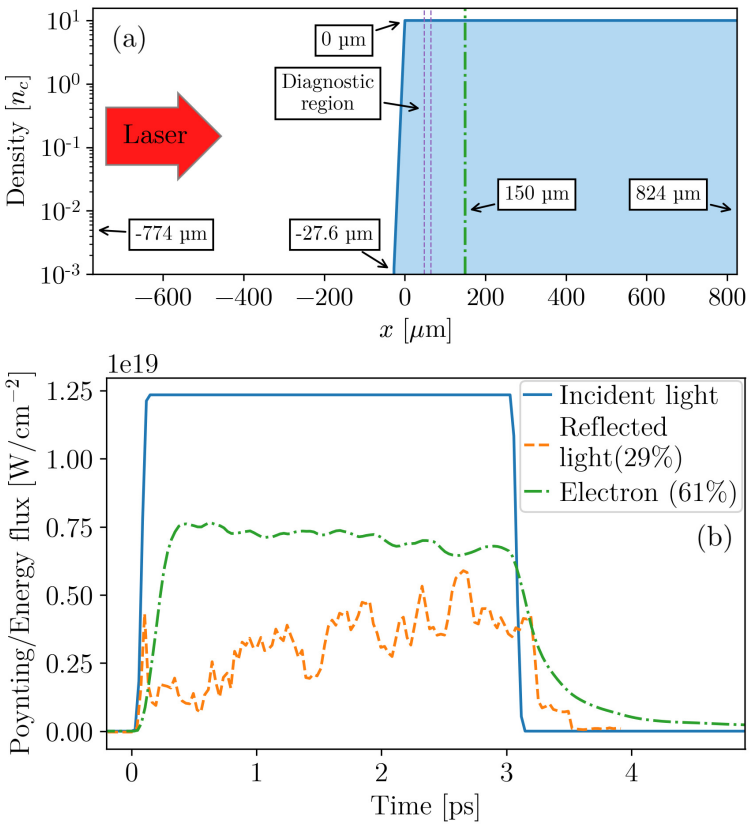

Figure A.3: (a) Simulation schematic, showing the full box size. The box is truncated at 150 μm when the absorber is in use. (b) Laser Poynting flux incident at the plasma critical interface, reflected Poynting flux measured 387 μm to the left of the critical interface and forward electron energy flux measured over the diagnostic region. All quantities are synced up in time for better visualization. Percentages represent integrated energy flux as a fraction of the total incident energy.

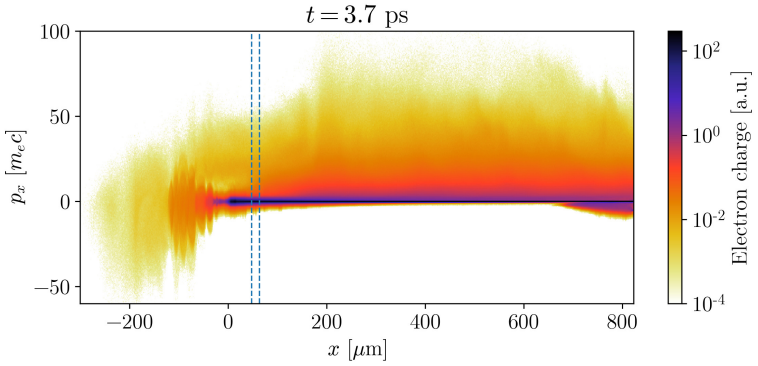

Figure A.4: The p_x-x phasespace for the causally separated simulation (single run, not averaged). The dashed lines indicate the diagnostic region, but the plasma has to be much larger in length to be causally separated from the hot return current reflecting off the right boundary.

different times with varied random number seeds and observed a factor of 2–4 variation in particle number in the tail of the momentum distribution over the diagnostic region. For this reason we performed the simulations presented in this appendix three times with different random number seeds. Unless otherwise noted, visualizations presented here are of data averaged over three different runs; this averaging gives increased confidence that any observed deviations from the causally separated run are due to the particle boundary conditions.

A.4.2 Effect of the absorber boundary condition

To greatly reduce computation time and resources, we desire to shrink the simulation region shown in Fig. A.4, but preserve the behavior from the causally separated run. We truncate the plasma at a distance of 150 μm (29025 cells in x) and vary the length of the absorber, where each absorber is designed to stop all hot particles 5 μm short of the right boundary. For all results shown here we use the linearly varying absorber from Sec. A.3.2 and calculate the local temperature via Eq. (A.11). We quote the mean free path for each absorber, which as shown in Fig. A.2 is 26% of the entire absorber length. We used an energy threshold of 6 times the local thermal velocity and re-emitted stopped particles at the local temperature.

Stopping was performed every time step for both electrons and ions to give a large number of stopping loops for a fast particle traversing the absorbing region. Particles are typically stopped every ~ 3 time steps, but we perform a stopping loop every time step to more accurately assess the different methods. Though it is much more important to use an absorber for electrons than for ions, we observed a sufficient number of hot ions reaching the thermal boundary to warrant stopping ions as well. Stopping loops were delayed until hot particles approached the absorber region. Particle recombination (for electrons) was executed every 5 time steps over the absorbing region; this dramatically reduces the simulation runtime as hot particles that have been split into 32 smaller particles are all stopped over a very short distance.

The p_x-x phasespaces for the causally separated ($\lambda = \infty$, where we are zooming in on a particular region), absorber (with $\lambda = 100\, c/\omega_p$) and no-absorber/truncated ($\lambda = 0$) simulations are shown in Figs. A.5 and A.6 at two different times. After just 1.1 ps, a hot reflux of electrons is visible in the truncated run [see Fig. A.5(c)] that has already entered the diagnostic region. These refluxing electrons are seen to completely overwhelm the simulation late in time [see Fig. A.6(c)], while the simulation with the absorber [see Fig. A.6(b)] is able to maintain an appropriate return current. These plots are not averaged over three simulations, so sizeable variations within the pre-plasma are expected for the causally separated run due to differences in random number initialization with a different box size [note that the phasespace in the density upramp and surrounding region are identical in Figs. A.5(b) and (c)].

To better visualize temporal behavior, we plot the electron energy flux in the x direction as a function of time and space for the causally separated, absorber, and no-absorber simulations in Fig. A.7. For the causally separated simulation, a steady stream of energy flux is observed to the right of the critical-density interface, which is slowly pushed forward in time. Energetic electrons are also seen to escape to the left as the plasma expands. This expansion is enhanced after the laser turns off. When using the absorber with $\lambda = 100\, c/\omega_p$, the energy flux looks qualitatively very similar to the causally separated run, except that the energy flux quickly

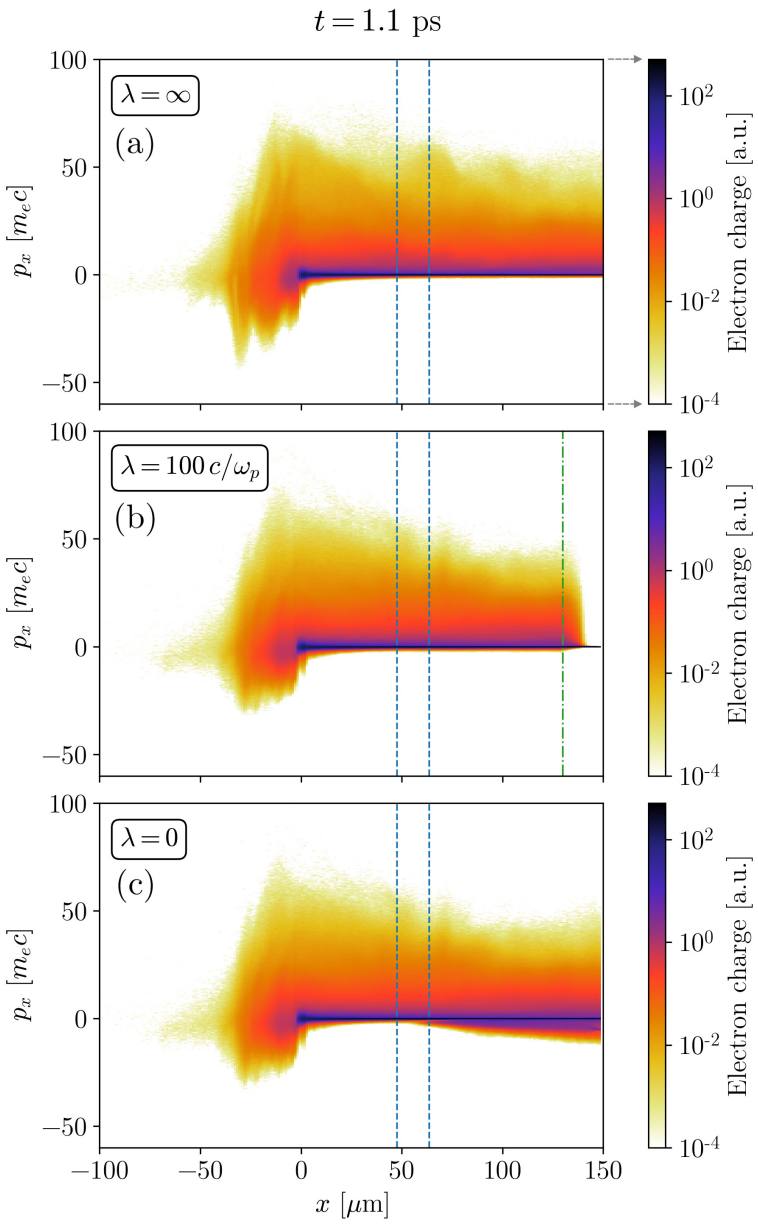

Figure A.5: The p_x-x phasespace (single runs, not averaged) for the (a) causally separated ($\lambda = \infty$), (b) absorber, and (c) no-absorber ($\lambda = 0$) simulations 1.1 ps after the incident laser. A hot reflux of electrons is already shown to be entering the dashed diagnostic region in (c) for the truncated run with no absorber.

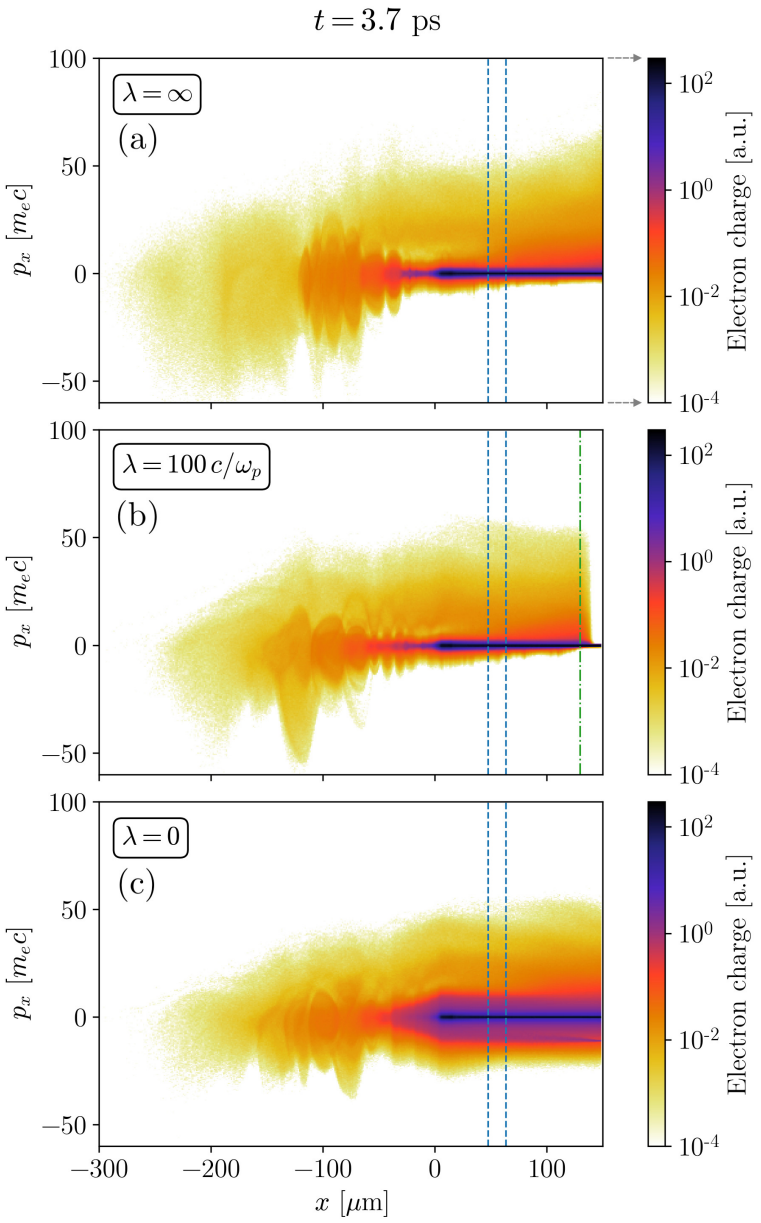

Figure A.6: The p_x-x phasespace (single runs, not averaged) for the (a) causally separated, (b) absorber, and (c) no-absorber simulations 3.7 ps after the incident laser. The refluxing electrons in (c) for the truncated run have completely altered the particle phasespace; the returning hot electrons cyclically interact with the laser and re-enter the plasma, artificially heating the bulk plasma to a much higher temperature than in the (a) casually separated or (b) absorbing runs.

decreases to zero in the absorber region. In contrast, the truncated simulation ($\lambda = 0$) shows that a large fraction of the forward energy flux is reflected from the right boundary (especially visible at 0.8 ps), so much so that it dramatically reduces the overall energy flux as it travels backward. Once the first reflux arrives back to the laser–plasma interface at around 1.5 ps, the forward energy flux is then permanently altered. This change in physics, as the hot return current interacts with and is accelerated by the laser, is the primary issue that the absorber is able to eliminate. Finally, this hot reflux of electrons is also visible in the blue negative energy flux after the laser turns off in the truncated run.

We also examine energy conservation (fields plus particles) across the simulation region when the absorber boundary condition is in use. To do this we compute the integral of energy density over a specific domain (V) and add the energy flux through the left and right boundaries of that domain (∂V):

$$\int_V U \, dV + \oint_{\partial V} \mathbf{S} \cdot d\mathbf{A}, \tag{A.9}$$

where U is the energy density $[E^2/8\pi + B^2/8\pi + \sum(\gamma - 1)m_e c^2]$ and $\mathbf{S}$ is the energy flux $[\mathbf{E} \times \mathbf{B}/4\pi + \sum(\gamma - 1)m_e c^2 \mathbf{p}/\gamma]$. We compute a running sum of this value over the simulation time (which should remain at zero) and then divide by the maximum energy present in the simulation box at any given time. This gives a good measure of the energy conservation of the code, although it is not perfect since we only use data reported every 401 time steps (0.3 ps). In Fig. A.8 we plot Eq. (A.9) as a function of time, where the right-hand side of volume V (i.e., the location of ∂V on the right) is given by the x coordinate displayed for an absorber with mean free path $\lambda = 100\, c/\omega_p$. We can see that to the left of the absorber (dashed line), the deviation in the coarsely computed energy conservation is less than 1.4%. However, by including the absorber region we see that a large fraction of the energy is steadily removed as energetic particles are stopped. Once again, it is this extended slowing of the particle beam that allows for an appropriate return current to develop, causing plasma to return back into the main simulation region.

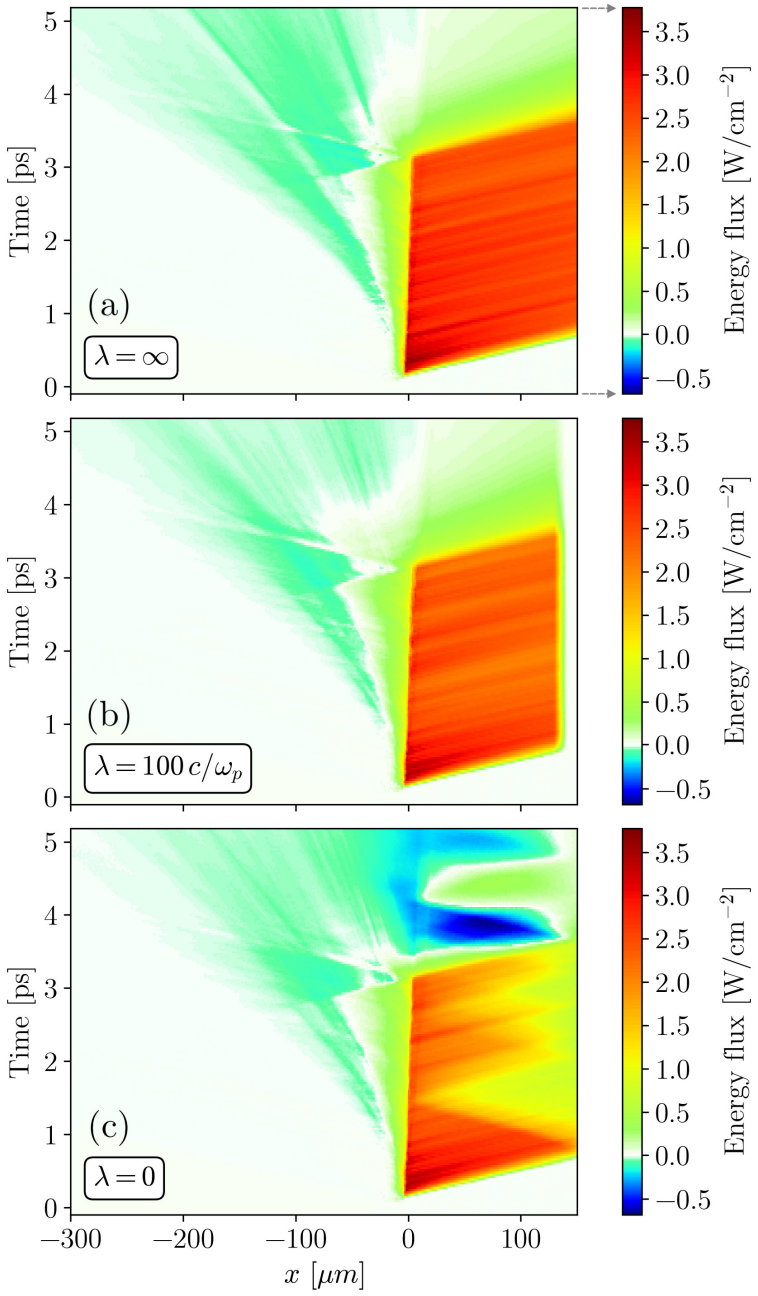

Figure A.7: Forward particle energy flux as a function of position and time for the (a) causally separated, (b) absorber, and (c) no-absorber simulations. For the truncated simulation in (c), the forward energy flux can be seen to be neutralized by a refluxing current emitted from the boundary. The absorber in (b) effectively reduces the particle energy flux before the simulation boundary without a reflux current.

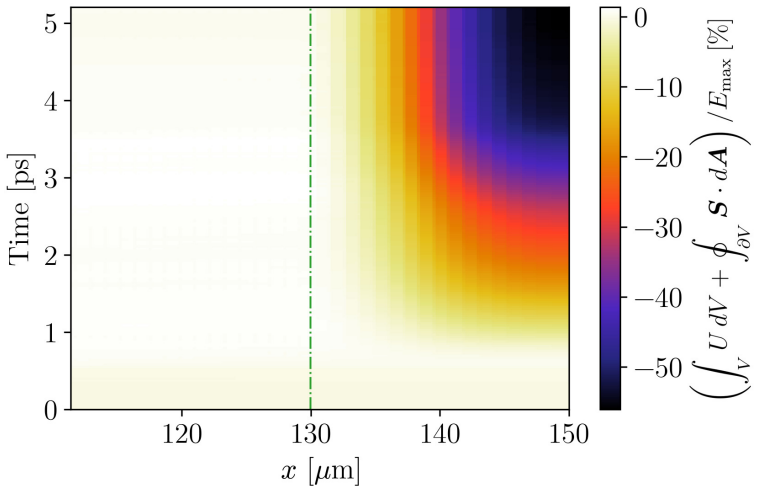

Figure A.8: The scaled deviation in energy conservation [see Eq. (A.9)] as a function of time, including all points to the left of a given x value (single run, not averaged). To the left of the absorbing region, energy is well conserved (<1.4% error), but in the absorbing region energy is steadily removed as particles are stopped.

A.4.3 Variation of absorber parameters

As mentioned in Sec. A.3 and Section A.6, there are a variety of options for implementing the absorber region. When determining the energy threshold and re-emission temperature of the stopped particles, we can calculate the background temperature dynamically by weighting the distribution function with the proper velocity to some power, or we can simply specify a constant value to use. Using a lower power (such as the fourth root) for the proper velocity will emphasize the bulk over a hot tail; more details are given in Section A.6. However, using the fourth-root temperature never improved the absorber performance for the simulations shown here, so we calculate the temperature in each cell as given by Eq. (A.11).

We can also use the hazard function probability defined in Sec. A.3.1 or the linearly varying probability defined in Sec. A.3.2 to stop the particles. In our tests these two choices produce similar results, but overall the linearly varying absorber maintained the proper response for a longer time. The main reason for this is that due to the periodicity in y, simulations

using the hazard-function absorber exhibited a large and increasing transverse temperature in the absorbing region; the hazard-function absorber preferentially stops particles with large forward momentum, allowing energetic particles to stream transversely and for some accelerating/reflecting fields to develop (see last paragraph of Sec. A.3.1). For this reason we use the linearly varying absorber in this appendix, which stops particles as a function of the magnitude of the velocity and not just the longitudinal component.

We compare a combination of absorbers in Fig. A.9, where we show the p_x momentum phasespace for all electrons in the diagnostic region at two different times. Although all absorber schemes appear to perform equally well early in time, the return current is clearly hotter when constant values of the energy threshold and re-emission temperature are given. For the static temperature simulation, we set the absorber to stop particles with energy greater than 0.6 keV and to re-emit particles at 0.1 keV (the original plasma temperature); in contrast the dynamic absorber stops particles moving at more than 6 times the locally computed thermal velocity. Using a static temperature performs poorly because, as seen even in the absorbing region of Fig. A.6(b), the plasma heats up significantly in response to the energetic electron beam. Particles stopped and re-emitted at the original temperature are not moving fast enough to provide the necessary return current, and a nonphysical potential develops that accelerates electrons backward with too much energy. Calculating the local temperature instead allows the absorber to accurately compensate for this dynamic behavior.

Although not shown here, we performed a series of simulations varying the mean free path of the absorber by factors of two between $\lambda = 0.1\, c/\omega_p$ and $\lambda = 200\, c/\omega_p$. We observed that if the absorber had a mean free path $\lambda \gtrsim 6\, c/\omega_p$, it was able to closely match the causally separated momentum distribution when averaged over three separate runs. However, individual simulations with $\lambda \lesssim 20\, c/\omega_p$ exhibited slightly greater variability in comparison to the causally separated data. In our simulations, an absorber with a mean free path of $6\, c/\omega_p$ performed only ~ 30 stopping events before nearly all particles were stopped, which was sufficient for a laser 3 ps in duration with $a_0 = 3$. However, care must be taken for lasers

of longer duration or higher intensity; Fig. A.9 shows that some absorbers can perform well (a) initially, but (b) eventually fail due to the large amount of energetic particles striking the absorber. Thus $\lambda \gtrsim \mathcal{O}(10\, c/\omega_p)$ gives a reasonable estimate of the appropriate mean free path, but the absorber length should be verified for each individual simulation.

A.4.4 Best practices

Here we make a few notes on best practices for performing simulations with the extended absorbing boundary condition. We found it important to also causally separate the vacuum boundary (where the laser is injected) from the laser–plasma interface. Even with absorbing particle boundary conditions at this vacuum boundary, most energetic particles that reached the vacuum boundary were immediately reflected back into the simulation space. This is a combined effect of the laser potential at the wall and the electric field buildup from exiting particles (a nonnegligible number of particles are accelerated toward the laser from the pre-plasma region). Refluxing from the vacuum boundary leads to a modified distribution at the laser–plasma interaction region, which then artificially inflates the forward electron energy flux in the target.

For simulations with a finite-width laser, absorber regions can also be placed at the transverse simulation edges to correctly handle the large flux of relativistic electrons expelled transversely from the laser spot. However, the effectiveness of the absorber relies on having a large number of particles in each cell (for calculating the temperature). If absorbers are placed at the transverse simulation boundaries, they may overlap with near-vacuum regions in and before the pre-plasma. Thus for finite-size-target simulations with multiple absorbers, we found it is useful to transition the absorbers positioned along the transverse boundaries to stop and re-emit particles based on a static (rather than dynamically calculated) temperature in those near-vacuum regions.

Finally, the start of the absorber region should be located a reasonable distance away from where accurate plasma measurements are expected. For example, when comparing

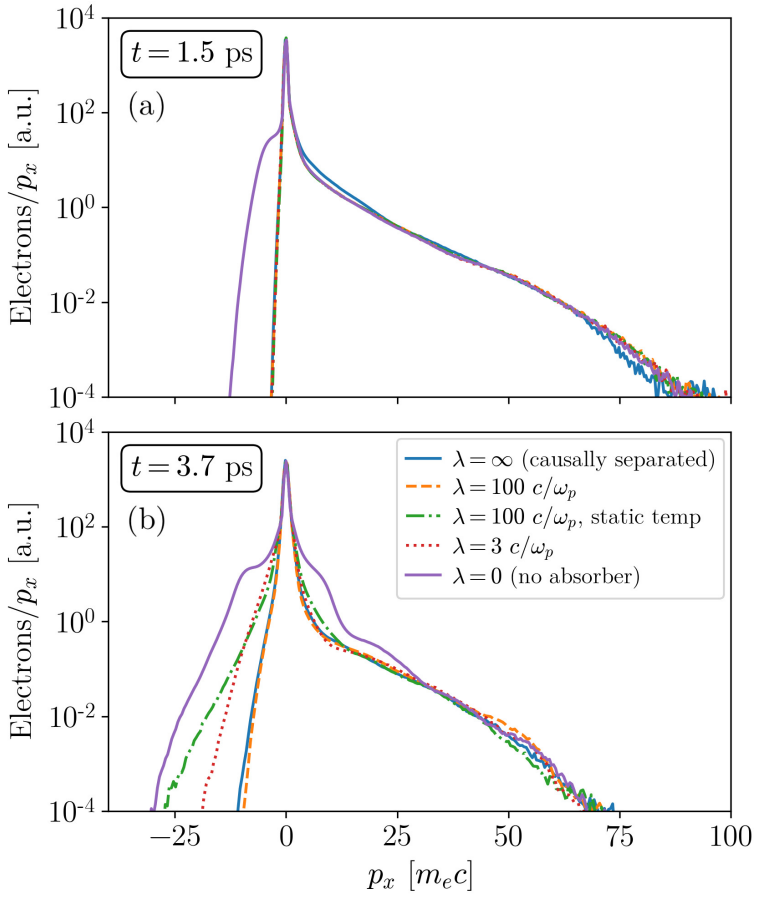

Figure A.9: The p_x phasespace for all electrons in the region 48–64 μm into the constant-density plasma for various schemes (a) 1.5 ps and (b) 3.7 ps after the laser was incident on the plasma. Though the performance of all shown absorbers is nearly identical early in time, either using a static temperature threshold and re-emission or using a very short absorber gives improper results later in time.

Figs. A.6(a) and A.6(b), the phasespace immediately in front of the absorbing region in (b) does not exactly mimic the causally separated phasespace in (a). Examining the particle phasespace for irregularities near the absorber region can help determine the appropriate distance at which to measure plasma quantities.

A.4.5 Future work

The implementation described here, though effective, is by no means a comprehensive treatment or unique solution to the reflux problem. Here we list some ideas that could be used to iterate on our proposed solution. Particles could be re-emitted from a distribution that is hotter in the return direction than in the forward direction, assisting in establishing the appropriate return current. Particles could be stopped preferentially based on their direction of motion. We employed absorbers for both ions and electrons in these simulations, but the ion response and stopping could be explored in greater detail for long-time simulations. Alternatives that are more computationally expensive could include applying a drag force to energetic particles over the length of the entire absorber [171, 209] (requires flagging individual particles) or calculating the re-emission temperature from a position located before the absorber region (requires global information). Last, it may be possible to develop a thermal bath boundary where particles are re-emitted from a non-Maxwellian distribution determined from a region somewhere near the boundary of the plasma.

A.5 Conclusion

Particle-in-cell simulations are useful for investigating intense laser–plasma interactions in overdense plasmas, but a truncated plasma boundary can produce an unphysically hot return current. This return current is present with absorbing, reflecting and thermal particle boundary conditions alike, and it can drastically alter simulation results. We have devised an absorbing particle boundary condition that stops energetic particles over a defined region of

the simulation space. Stopping these particles over a sufficiently large distance allows the background plasma to generate a suitably cool return current that mimics the results of a semi-infinite, causally separated simulation.

Various different schemes were proposed for statistically selecting, stopping and re-emitting hot particles, with the best results given by the linearly varying absorber described in Sec. A.3.2 that calculates the local temperature via Eq. (A.11). The appropriate mean free path of the absorber was explored, showing that an absorber with a mean free path of $\lambda \gtrsim \mathcal{O}(10\, c/\omega_p)$ gives proper results for our tests. As simulation behavior can vary greatly depending on the application, care must be taken to ensure that the absorber parameters used for a particular case appropriately mimic the behavior of a semi-infinite boundary.

A.6 Supplemental—Computing local temperature

In this supplemental section we describe how we calculate the local proper thermal velocity for use with the damping parameters. First, we assume a Maxwellian distribution of the form

$$f_0(u) = \frac{n_0}{\sqrt{2\pi\bar{u}_0^2}} \exp\left(\frac{u^2}{2\bar{u}_0^2}\right) \tag{A.10}$$

for density n_0, proper velocity $u \equiv \gamma v$ and thermal velocity $\bar{u}_0$. We can approximate the average thermal velocity by summing $|u|$ over all particles in a cell. For a Maxwellian of the same form as in Eq. (A.10), we have that

$$\int_{-\infty}^{\infty} |u| f_0(u)\, du = \sqrt{\frac{2}{\pi}} n_0 \bar{u}_0, \tag{A.11}$$

which can be inverted to find $\bar{u}_0$.

If instead the plasma is considered to be two distinct species with different thermal velocities and densities [e.g., including a beam with density n_1, thermal velocity $\bar{u}_1$ and corresponding distribution $f_1(u)$], the above integral in Eq. (A.11) could yield a distorted thermal velocity. Another option is to perform the integral using the fourth-root of the

thermal velocity, which gives

$$\int_{-\infty}^{\infty} |u|^{1/4} \left[f_0(u) + f_1(u) \right] \, du = \frac{2^{1/8} \Gamma\left(\frac{5}{8}\right)}{\sqrt{\pi}} \left(n_0 \bar{u}_0^{1/4} + n_1 \bar{u}_1^{1/4} \right), \tag{A.12}$$

where Γ is the standard gamma function. Since the densities add linearly but the thermal velocities add as the fourth-root, a high-energy beam should distort the sum less than in Eq. (A.11). Note that Eq. (A.12) should be calculated for a single distribution and inverted to solve for $\bar{u}_0$. In practice, we found that the fourth-root calculation differed only slightly from the simple average of the thermal velocity due to the extremely low density of the beam. However, the fourth-root calculation may be important for some parameter regimes.

REFERENCES

[1] J L Shaw, N Lemos, L D Amorim, N Vafaei-Najafabadi, K A Marsh, F S Tsung, W B Mori, and C Joshi. Role of Direct Laser Acceleration of Electrons in a Laser Wakefield Accelerator with Ionization Injection. *Physical Review Letters*, 118(6):64801, 2017.

[2] Gerard A. Mourou, Toshiki Tajima, and Sergei V. Bulanov. Optics in the relativistic regime. *Reviews of Modern Physics*, 78(2):309–371, apr 2006.

[3] Gerard Mourou and Toshiki Tajima. More Intense, Shorter Pulses. *Science*, 331(6013):41–42, jan 2011.

[4] Colin Danson, David Hillier, Nicholas Hopps, and David Neely. Petawatt class lasers worldwide. *High Power Laser Science and Engineering*, 3:e3, jan 2015.

[5] John Lindl. Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain. *Physics of Plasmas*, 2(11):3933–4024, nov 1995.

[6] R. P. Drake, J. J. Carroll, T. B. Smith, P. Keiter, S. Gail Glendinning, Omar Hurricane, Kent Estabrook, D. D. Ryutov, B. A. Remington, R J Wallace, Eli Michael, and R. McCray. Laser experiments to simulate supernova remnants. *Physics of Plasmas*, 7(5):2142–2148, may 2000.

[7] A. B. Reighard, R. P. Drake, K. K. Dannenberg, D. J. Kremer, M. Grosskopf, E. C. Harding, D. R. Leibrandt, S. G. Glendinning, T. S. Perry, B. A. Remington, J. Greenough, J. Knauer, T. Boehly, S. Bouquet, L. Boireau, M. Koenig, and T. Vinci. Observation of collapsing radiative shocks in laboratory experiments. *Physics of Plasmas*, 13(8):082901, aug 2006.

[8] C. Joshi, B. Blue, C. E. Clayton, E. Dodd, C. Huang, K. A. Marsh, W. B. Mori, S. Wang, M. J. Hogan, C. O'Connell, R. Siemann, D. Watz, P. Muggli, T. Katsouleas, and S. Lee. High energy density plasma science with an ultrarelativistic electron beam. *Physics of Plasmas*, 9(5):1845–1855, may 2002.

[9] C. Joshi, S. Corde, and W. B. Mori. Perspectives on the generation of electron beams from plasma-based accelerators and their near and long term applications. *Physics of Plasmas*, 27(7):070602, jul 2020.

[10] M. D. Perry, J. A. Sefcik, T. Cowan, S. Hatchett, A. Hunt, M. Moran, D. Pennington, R. Snavely, and S. C. Wilks. Hard x-ray production from high intensity laser solid interactions (invited). *Review of Scientific Instruments*, 70(1):265–269, jan 1999.

[11] K. K. Kainz, K. R. Hogstrom, J. A. Antolak, P. R. Almond, C. D. Bloch, C. Chiu, M. Fomytskyi, F. Raischel, M. Downer, and T. Tajima. Dose properties of a laser

accelerated electron beam and prospects for clinical application. *Medical Physics*, 31(7):2053–2067, jun 2004.

[12] C. M. Brenner, S. R. Mirfayzi, D. R. Rusby, C. Armstrong, A. Alejo, L. A. Wilson, R. Clarke, H. Ahmed, N M H Butler, D. Haddock, A. Higginson, A. McClymont, C. Murphy, M. Notley, P. Oliver, R. Allott, C. Hernandez-Gomez, S. Kar, P. McKenna, and D. Neely. Laser-driven x-ray and neutron source development for industrial applications of plasma accelerators. *Plasma Physics and Controlled Fusion*, 58(1):014039, jan 2016.

[13] Christopher P. Jones, Ceri M. Brenner, Camilla A. Stitt, Chris Armstrong, Dean R. Rusby, Seyed R. Mirfayzi, Lucy A. Wilson, Aarón Alejo, Hamad Ahmed, Ric Allott, Nicholas M.H. Butler, Robert J. Clarke, David Haddock, Cristina Hernandez-Gomez, Adam Higginson, Christopher Murphy, Margaret Notley, Charilaos Paraskevoulakos, John Jowsey, Paul McKenna, David Neely, Satya Kar, and Thomas B. Scott. Evaluating laser-driven Bremsstrahlung radiation sources for imaging and analysis of nuclear waste packages. *Journal of Hazardous Materials*, 318:694–701, nov 2016.

[14] C. Courtois, A. Compant La Fontaine, O. Landoas, G. Lidove, V. Méot, P. Morel, R. Nuter, E. Lefebvre, A. Boscheron, J. Grenier, M. M. Aléonard, M. Gerbaux, F. Gobet, F. Hannachi, G. Malka, J. N. Scheurer, and M. Tarisien. Effect of plasma density scale length on the properties of bremsstrahlung x-ray sources created by picosecond laser pulses. *Physics of Plasmas*, 16(1):013105, jan 2009.

[15] D. Babonneau, M. Primout, F. Girard, J.-P. Jadaud, M. Naudy, B. Villette, S. Depierreux, C. Blancard, G. Faussurier, K. B. Fournier, L. Suter, R. Kauffman, S. Glenzer, M. C. Miller, J. Grün, and J. Davis. Efficient multi-keV X-ray sources from laser-exploded metallic thin foils. *Physics of Plasmas*, 15(9):092702, sep 2008.

[16] M.A. Barrios, K.B. Fournier, S.P. Regan, O. Landen, M. May, Y.P. Opachich, K. Widmann, D.K. Bradley, and G.W. Collins. Backlighter development at the National Ignition Facility (NIF): Zinc to zirconium. *High Energy Density Physics*, 9(3):626–634, sep 2013.

[17] D. R. Rusby, C. M. Brenner, C. Armstrong, L. A. Wilson, R. Clarke, A. Alejo, H. Ahmed, N. M. H. Butler, D. Haddock, A. Higginson, A. McClymont, S. R. Mirfayzi, C. Murphy, M. Notley, P. Oliver, R. Allott, C. Hernandez-Gomez, S. Kar, P. McKenna, and D. Neely. Pulsed x-ray imaging of high-density objects using a ten picosecond high-intensity laser driver. In Keith L Lewis and Richard C Hollins, editors, *Emerging Imaging and Sensing Technologies*, volume 9992, pages 61–68. International Society for Optics and Photonics, SPIE, oct 2016.

[18] C. Courtois, R. Edwards, A. Compant La Fontaine, C. Aedy, M. Barbotin, S. Bazzoli, L. Biddle, D. Brebion, J. L. Bourgade, D. Drew, M. Fox, M. Gardner, J. Gazave, J. M.

Lagrange, O. Landoas, L. Le Dain, E. Lefebvre, D. Mastrosimone, N. Pichoff, G. Pien, M. Ramsay, A. Simons, N. Sircombe, C. Stoeckl, and K. Thorp. High-resolution multi-MeV x-ray radiography using relativistic laser-solid interaction. *Physics of Plasmas*, 18(2):023101, feb 2011.

[19] Y. Glinec, J. Faure, L. Le Dain, S. Darbon, T. Hosokai, J. J. Santos, E. Lefebvre, J. P. Rousseau, F. Burgy, B. Mercier, and V. Malka. High-Resolution γ-Ray Radiography Produced by a Laser-Plasma Driven Electron Source. *Physical Review Letters*, 94(2):025003, jan 2005.

[20] R. D. Edwards, M. A. Sinclair, T. J. Goldsack, K. Krushelnick, F. N. Beg, E. L. Clark, A. E. Dangor, Z. Najmudin, M. Tatarakis, B. Walton, M. Zepf, K. W. D. Ledingham, I. Spencer, P. A. Norreys, R. J. Clarke, R. Kodama, Y. Toyama, and M. Tampo. Characterization of a gamma-ray source based on a laser-plasma accelerator with applications to radiography. *Applied Physics Letters*, 80(12):2129–2131, mar 2002.

[21] W. P. Leemans, D. Rodgers, P. E. Catravas, C. G. R. Geddes, G. Fubiani, E. Esarey, B. A. Shadwick, R. Donahue, and A. Smith. Gamma-neutron activation experiments using laser wakefield accelerators. *Physics of Plasmas*, 8(5):2510–2516, may 2001.

[22] John M Dawson. Particle simulation of plasmas. *Reviews of Modern Physics*, 55(2):403–447, apr 1983.

[23] C.K. Birdsall and A.B Langdon. *Plasma Physics Via Computer Simulation*. McGraw Hill Higher Education, oct 1984.

[24] Kane S Yee. Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, may 1966.

[25] Jay P Boris and Ramy A Shanny. *Proceedings: Fourth Conference on Numerical Simulation of Plasmas, November 2, 3, 1970*. Naval Research Laboratory, 1972.

[26] J P Boris. Relativistic plasma simulation—Optimization of a hybrid code. In *4th Conference on Numerical Simulation of Plasmas*, pages 3–67, Washington D. C., 1970. Naval Research Laboratory.

[27] A Bruce Langdon. On enforcing Gauss' law in electromagnetic particle-in-cell codes. *Computer Physics Communications*, 70(3):447–450, 1992.

[28] R A Fonseca, L O Silva, F S Tsung, V K Decyk, W Lu, C Ren, W B Mori, S Deng, S Lee, T Katsouleas, and J C Adam. OSIRIS: A Three-Dimensional, Fully Relativistic Particle in Cell Code for Modeling Plasma Based Accelerators. In Peter M A Sloot, Alfons G Hoekstra, C J Kenneth Tan, and Jack J Dongarra, editors, *Computational Science — ICCS 2002: International Conference Amsterdam, The Netherlands, April*

21–24, 2002 Proceedings, Part III, pages 342–351. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.

[29] Roy G Hemker. *Particle-in-cell modeling of plasma-based accelerators in two and three dimensions*. Phd, University of California, Los Angeles, 2000.

[30] A. Davidson, A. Tableman, Weiming An, F.S. Tsung, Wei Lu, Jorge Vieira, R.A. Fonseca, L.O. Silva, and W.B. Mori. Implementation of a hybrid particle code with a PIC description in r–z and a gridless description in ϕ into OSIRIS. *Journal of Computational Physics*, 281:1063–1077, jan 2015.

[31] T.Zh. Zh Esirkepov. Exact charge conservation scheme for Particle-in-Cell simulation with an arbitrary form-factor. *Computer Physics Communications*, 135(2):144–153, apr 2001.

[32] J.-L. Vay. Simulation of beams or plasmas crossing at relativistic velocity. *Physics of Plasmas*, 15(5):056701, may 2008.

[33] A. V. Higuera and J. R. Cary. Structure-preserving second-order integration of relativistic charged particle trajectories in electromagnetic fields. *Physics of Plasmas*, 24(5):052104, may 2017.

[34] Fei Li, Viktor K. Decyk, Kyle G. Miller, Adam Tableman, Frank S. Tsung, Marija Vranic, Ricardo A. Fonseca, and Warren B. Mori. Accurately simulating nine-dimensional phase space of relativistic particles in strong fields. *Journal of Computational Physics*, 438:110367, aug 2021.

[35] Brendan B. Godfrey, Jean-Luc Vay, and Irving Haber. Numerical stability analysis of the pseudo-spectral analytical time-domain PIC algorithm. *Journal of Computational Physics*, 258:689–704, feb 2014.

[36] Xinlu Xu, Fei Li, Frank S. Tsung, Thamine N. Dalichaouch, Weiming An, Han Wen, Viktor K. Decyk, Ricardo A. Fonseca, Mark J. Hogan, and Warren B. Mori. On numerical errors to the fields surrounding a relativistically moving particle in PIC codes. *Journal of Computational Physics*, 413:109451, jul 2020.

[37] Fei Li, Kyle G. Miller, Xinlu Xu, Frank S. Tsung, Viktor K. Decyk, Weiming An, Ricardo A. Fonseca, and Warren B. Mori. A new field solver for modeling of relativistic particle-laser interactions using the particle-in-cell algorithm. *Computer Physics Communications*, 258:107580, jan 2021.

[38] E.L. Lindman. "Free-space" boundary conditions for the time dependent wave equation. *Journal of Computational Physics*, 18(1):66–78, may 1975.

[39] Jean-Luc Vay. Asymmetric Perfectly Matched Layer for the Absorption of Waves. *Journal of Computational Physics*, 183(2):367–399, dec 2002.

[40] Paulett C Liewer and Viktor K Decyk. A general concurrent algorithm for plasma particle-in-cell simulation codes. *Journal of Computational Physics*, 85(2):302–322, dec 1989.

[41] Viktor K Decyk and Tajendra V Singh. Particle-in-Cell algorithms for emerging computer architectures. *Computer Physics Communications*, 185(3):708–719, 2014.

[42] Kai Germaschewski, William Fox, Stephen Abbott, Narges Ahmadi, Kristofor Maynard, Liang Wang, Hartmut Ruhl, and Amitava Bhattacharjee. The Plasma Simulation Code: A modern particle-in-cell code with patch-based load-balancing. *Journal of Computational Physics*, 318:305–326, 2016.

[43] R A Fonseca, J Vieira, F Fiuza, A Davidson, F S Tsung, W B Mori, and L O Silva. Exploiting multi-scale parallelism for large scale numerical modelling of laser wakefield accelerators. *Plasma Physics and Controlled Fusion*, 55(12):124011, dec 2013.

[44] K. M. Schoeffler, T. Grismayer, D. Uzdensky, R. A. Fonseca, and L. O. Silva. Bright Gamma-Ray Flares Powered by Magnetic Reconnection in QED-strength Magnetic Fields. *The Astrophysical Journal*, 870(1):49, 2019.

[45] Peicheng Yu, Xinlu Xu, Viktor K Decyk, Weiming An, Jorge Vieira, Frank S Tsung, Ricardo A Fonseca, Wei Lu, Luis O Silva, and Warren B Mori. Modeling of laser wakefield acceleration in Lorentz boosted frame using EM-PIC code with spectral solver. *Journal of Computational Physics*, 266:124–138, 2014.

[46] Peicheng Yu. *Lorentz boosted frame simulation technique in Particle-in-cell methods*. PhD thesis, University of California, Los Angeles, 2016.

[47] Peicheng Yu, Xinlu Xu, Asher Davidson, Adam Tableman, Thamine Dalichaouch, Fei Li, Michael D. Meyers, Weiming An, Frank S. Tsung, Viktor K. Decyk, Frederico Fiuza, Jorge Vieira, Ricardo A. Fonseca, Wei Lu, Luis O. Silva, and Warren B. Mori. Enabling Lorentz boosted frame particle-in-cell simulations of laser wakefield acceleration in quasi-3D geometry. *Journal of Computational Physics*, 316:747–759, jul 2016.

[48] L Yin, B J Albright, H A Rose, D S Montgomery, J L Kline, R K Kirkwood, P Michel, K J Bowers, and B Bergen. Self-organized coherent bursts of stimulated Raman scattering and speckle interaction in multi-speckled laser beams. *Physics of Plasmas*, 20(1):12702, 2013.

[49] W.B. Mori. The physics of the nonlinear optics of plasmas at relativistic intensities for short-pulse lasers. *IEEE Journal of Quantum Electronics*, 33(11):1942–1953, 1997.

[50] John M. Dawson. Nonlinear Electron Oscillations in a Cold Plasma. *Physical Review*, 113(2):383–387, jan 1959.

[51] A. Akhiezer and R. Polovin. Theory of Wave Motion of an Electron Plasma. *Soviet Phys. JETP*, Vol: 3:696–705, 1956.

[52] T Tajima and J M Dawson. Laser electron accelerator. *Physical Review Letters*, 43(4):267–270, jul 1979.

[53] Juwen Wang. *RF properties of periodic accelerating structures for linear colliders*. Phd, Stanford University, 1989.

[54] K. Nordlund and F. Djurabekova. Defect model for the dependence of breakdown rate on external electric fields. *Physical Review Special Topics - Accelerators and Beams*, 15(7):071002, jul 2012.

[55] C. Joshi, T. Tajima, J. M. Dawson, H. A. Baldis, and N. A. Ebrahim. Forward Raman Instability and Electron Acceleration. *Physical Review Letters*, 47(18):1285–1288, nov 1981.

[56] C. Joshi, W. B. Mori, T. Katsouleas, J. M. Dawson, J. M. Kindel, and D. W. Forslund. Ultrahigh gradient particle acceleration by intense laser-driven plasma density waves. *Nature*, 311(5986):525–529, oct 1984.

[57] J Krall, A Ting, E Esarey, and P. Sprangle. Enhanced acceleration in a self-modulated-laser wake-field accelerator. *Physical Review E*, 48(3):2157–2161, sep 1993.

[58] A. Modena, Z. Najmudin, A. E. Dangor, C. E. Clayton, K. A. Marsh, C. Joshi, V. Malka, C. B. Darrow, C. Danson, D. Neely, and F. N. Walsh. Electron acceleration from the breaking of relativistic plasma waves. *Nature*, 377(6550):606–608, oct 1995.

[59] E. Esarey, P. Sprangle, J. Krall, and A. Ting. Overview of plasma-based accelerator concepts. *IEEE Transactions on Plasma Science*, 24(2):252–288, apr 1996.

[60] 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. *Physical Review Letters*, 54(7):693–696, feb 1985.

[61] B. E. Blue, C. E. Clayton, C. L. O'Connell, F.-J. Decker, M. J. Hogan, C. Huang, R. Iverson, C. Joshi, T. C. Katsouleas, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, R. Siemann, and D. Walz. Plasma-Wakefield Acceleration of an Intense Positron Beam. *Physical Review Letters*, 90(21):214801, may 2003.

[62] L.M. Gorbunov and V.I. Kirsanov. Excitation of plasma waves by an electromagnetic wave packet. *Soviet physics, JETP*, 66(2):290–294, 1987.

[63] N. E. Andreev, L. M. Gorbunov, V. I. Kirsanov, K. Nakajima, and A. Ogata. Structure of the wake field in plasma channels. *Physics of Plasmas*, 4(4):1145–1153, apr 1997.

[64] Patrick Mora and Thomas M. Antonsen. Electron cavitation and acceleration in the wake of an ultraintense, self-focused laser pulse. *Physical Review E*, 53(3):R2068–R2071, mar 1996.

[65] A. Pukhov and J. Meyer-ter Vehn. Laser wake field acceleration: the highly non-linear broken-wave regime. *Applied Physics B: Lasers and Optics*, 74(4-5):355–361, apr 2002.

[66] 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. *Physics of Plasmas*, 13(5):056709, may 2006.

[67] W. Lu, M. Tzoufras, C. Joshi, F. S. Tsung, W. B. Mori, J. Vieira, R. A. Fonseca, and L. O. Silva. Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime. *Physical Review Special Topics - Accelerators and Beams*, 10(6):061301, jun 2007.

[68] Alexander Pukhov. Strong field interaction of laser radiation. *Reports on Progress in Physics*, 66(1):47–101, jan 2003.

[69] Silvia Cipiccia, Mohammad R. Islam, Bernhard Ersfeld, Richard P. Shanks, Enrico Brunetti, Gregory Vieux, Xue Yang, Riju C. Issac, Samuel M. Wiggins, Gregor H. Welsh, Maria-Pia Anania, Dzmitry Maneuski, Rachel Montgomery, Gary Smith, Matthias Hoek, David J. Hamilton, Nuno R. C. Lemos, Dan Symes, Pattathil P. Rajeev, Val O. Shea, João M. Dias, and Dino A. Jaroszynski. Gamma-rays from harmonically resonant betatron oscillations in a plasma wake. *Nature Physics*, 7(11):867–871, nov 2011.

[70] C. Gahn, G. D. Tsakiris, A. Pukhov, J. Meyer-ter Vehn, G. Pretzler, P. Thirolf, D. Habs, and K. J. Witte. Multi-MeV Electron Beam Generation by Direct Laser Acceleration in High-Density Plasma Channels. *Physical Review Letters*, 83(23):4772–4775, dec 1999.

[71] Z. Najmudin, K. Krushelnick, E. L. Clark, S. P. D. Mangles, B. Walton, A. E. Dangor, S. Fritzler, V. Malka, E. Lefebvre, D. Gordon, F. S. Tsung, and C. Joshi. Self-modulated wakefield and forced laser wakefield acceleration of electrons. *Physics of Plasmas*, 10(5):2071–2077, may 2003.

[72] S. P. D. Mangles, B. R. Walton, M. Tzoufras, Z. Najmudin, R. J. Clarke, A. E. Dangor, R. G. Evans, S. Fritzler, A. Gopal, C. Hernandez-Gomez, W. B. Mori, W. Rozmus, M. Tatarakis, A. G. R. Thomas, F. S. Tsung, M. S. Wei, and K. Krushelnick. Electron Acceleration in Cavitated Channels Formed by a Petawatt Laser in Low-Density Plasma. *Physical Review Letters*, 94(24):245001, jun 2005.

[73] N Lemos, J L Martins, F S Tsung, J L Shaw, K A Marsh, F Albert, B B Pollock, and C Joshi. Self-modulated laser wakefield accelerators as x-ray sources. *Plasma Physics and Controlled Fusion*, 58(3):034018, mar 2016.

[74] P. M. King, K. Miller, N. Lemos, J. L. Shaw, B. F. Kraus, M. Thibodeau, B. M. Hegelich, J. Hinojosa, P. Michel, C. Joshi, K. A. Marsh, W. Mori, A. Pak, A. G. R. Thomas, and F. Albert. Predominant contribution of direct laser acceleration to high-energy electron spectra in a low-density self-modulated laser wakefield accelerator. *Physical Review Accelerators and Beams*, 24(1):011302, jan 2021.

[75] E. Esarey, B. A. Shadwick, P. Catravas, and W. P. Leemans. Synchrotron radiation from electron beams in plasma-focusing channels. *Physical Review E*, 65(5):056505, may 2002.

[76] F. Albert, B. B. Pollock, J. L. Shaw, K. A. Marsh, J. E. Ralph, Y.-H. Chen, D. Alessi, A. Pak, C. E. Clayton, S. H. Glenzer, and C. Joshi. Angular Dependence of Betatron X-Ray Spectra from a Laser-Wakefield Accelerator. *Physical Review Letters*, 111(23):235004, dec 2013.

[77] T. W. Huang, A. P.L. Robinson, C. T. Zhou, B. Qiao, B. Liu, S. C. Ruan, X. T. He, and P. A. Norreys. Characteristics of betatron radiation from direct-laser-accelerated electrons. *Physical Review E*, 93(6):063203, 2016.

[78] A. Pukhov, Z.-M. Sheng, and J. Meyer-ter Vehn. Particle acceleration in relativistic laser channels. *Physics of Plasmas*, 6(7):2847–2854, jul 1999.

[79] Xi Zhang, Vladimir N. Khudik, and Gennady Shvets. Synergistic Laser-Wakefield and Direct-Laser Acceleration in the Plasma-Bubble Regime. *Physical Review Letters*, 114(18):184801, may 2015.

[80] Xi Zhang, Vladimir N. Khudik, Alexander Pukhov, and Gennady Shvets. Laser wakefield and direct acceleration with ionization injection. *Plasma Physics and Controlled Fusion*, 58(3):034011, mar 2016.

[81] S. Kneip, S. R. Nagel, C. Bellei, N. Bourgeois, A. E. Dangor, A. Gopal, R. Heathcote, S. P. D. Mangles, J. R. Marquès, A. Maksimchuk, P. M. Nilson, K. Ta Phuoc, S. Reed, M. Tzoufras, F. S. Tsung, L. Willingale, W. B. Mori, A. Rousse, K. Krushelnick, and Z. Najmudin. Observation of Synchrotron Radiation from Electrons Accelerated in a Petawatt-Laser-Generated Plasma Cavity. *Physical Review Letters*, 100(10):105006, mar 2008.

[82] E. Esarey, C. B. Schroeder, and W. P. Leemans. Physics of laser-driven plasma-based electron accelerators. *Reviews of Modern Physics*, 81(3):1229–1285, aug 2009.

[83] F. Grüner, S. Becker, U. Schramm, T. Eichner, M. Fuchs, R. Weingartner, D. Habs, J. Meyer-ter Vehn, M. Geissler, M. Ferrario, L. Serafini, B. van der Geer, H. Backe, W. Lauth, and S. Reiche. Design considerations for table-top, laser-based VUV and X-ray free electron lasers. *Applied Physics B*, 86(3):431–435, feb 2007.

[84] Kazuhisa Nakajima. Towards a table-top free-electron laser. *Nature Physics*, 4(2):92–93, feb 2008.

[85] S. Corde, K. Ta Phuoc, G. Lambert, R. Fitour, V. Malka, A. Rousse, A. Beck, and E. Lefebvre. Femtosecond x rays from laser-plasma accelerators. *Reviews of Modern Physics*, 85(1):1–48, jan 2013.

[86] M. E. Couprie, M. Labat, C. Evain, F. Marteau, F. Briquez, M. Khojoyan, C. Benabderrahmane, L. Chapuis, N. Hubert, C. Bourassin-Bouchet, M El Ajjouri, F. Bouvet, Y. Dietrich, M. Valléau, G. Sharma, W. Yang, O. Marcouillé, J. Vétéran, P. Berteaud, T. El Ajjouri, L. Cassinari, C. Thaury, G. Lambert, I. Andriyash, V. Malka, X. Davoine, M. A. Tordeux, C. Miron, D. Zerbib, K. Tavakoli, J. L. Marlats, M. Tilmont, P. Rommeluère, J. P. Duval, M H N'Guyen, A. Rouqier, M. Vanderbergue, C. Herbeaux, M. Sebdouai, A. Lestrade, N. Leclercq, D. Dennetière, M. Thomasset, F. Polack, S. Bielawski, C. Szwaj, and A. Loulergue. An application of laser–plasma acceleration: towards a free-electron laser amplification. *Plasma Physics and Controlled Fusion*, 58(3):034020, mar 2016.

[87] Weiming An, Wei Lu, Chengkun Huang, Xinlu Xu, Mark J. Hogan, Chan Joshi, and Warren B. Mori. Ion Motion Induced Emittance Growth of Matched Electron Beams in Plasma Wakefields. *Physical Review Letters*, 118(24):244801, jun 2017.

[88] G. G. Manahan, A. F. Habib, P. Scherkl, P. Delinikolas, A. Beaton, A. Knetsch, O. Karger, G. Wittig, T. Heinemann, Z. M. Sheng, J. R. Cary, D. L. Bruhwiler, J. B. Rosenzweig, and B. Hidding. Single-stage plasma-based correlated energy spread compensation for ultrahigh 6D brightness electron beams. *Nature Communications*, 8(1):15705, aug 2017.

[89] T. André, I. A. Andriyash, A. Loulergue, M. Labat, E. Roussel, A. Ghaith, M. Khojoyan, C. Thaury, M. Valléau, F. Briquez, F. Marteau, K. Tavakoli, P. N'Gotta, Y. Dietrich, G. Lambert, V. Malka, C. Benabderrahmane, J. Vétéran, L. Chapuis, T. El Ajjouri, M. Sebdaoui, N. Hubert, O. Marcouillé, P. Berteaud, N. Leclercq, M. El Ajjouri, P. Rommeluère, F. Bouvet, J. P. Duval, C. Kitegi, F. Blache, B. Mahieu, S. Corde, J. Gautier, K. Ta Phuoc, J. P. Goddet, A. Lestrade, C. Herbeaux, C. Évain, C. Szwaj, S. Bielawski, A. Tafzi, P. Rousseau, S. Smartsev, F. Polack, D. Dennetière, C. Bourassin-Bouchet, C. De Oliveira, and M.-E. Couprie. Control of laser plasma accelerated electrons for light sources. *Nature Communications*, 9(1):1334, dec 2018.

[90] A. Ghaith, D. Oumbarek, E. Roussel, S. Corde, M. Labat, T. André, A. Loulergue, I. A. Andriyash, O. Chubar, O. Kononenko, S. Smartsev, O. Marcouillé, C. Kitégi, F. Marteau, M. Valléau, C. Thaury, J. Gautier, S. Sebban, A. Tafzi, F. Blache, F. Briquez, K. Tavakoli, A. Carcy, F. Bouvet, Y. Dietrich, G. Lambert, N. Hubert, M. El Ajjouri, F. Polack, D. Dennetière, N. Leclercq, P. Rommeluère, J.-P. Duval, M. Sebdaoui, C. Bourgoin, A. Lestrade, C. Benabderrahmane, J. Vétéran, P. Berteaud,

C. De Oliveira, J. P. Goddet, C. Herbeaux, C. Szwaj, S. Bielawski, V. Malka, and M.-E. Couprie. Tunable High Spatio-Spectral Purity Undulator Radiation from a Transported Laser Plasma Accelerated Electron Beam. *Scientific Reports*, 9(1):19020, dec 2019.

[91] K. Khrennikov, J. Wenz, A. Buck, J. Xu, M. Heigoldt, L. Veisz, and S. Karsch. Tunable All-Optical Quasimonochromatic Thomson X-Ray Source in the Nonlinear Regime. *Physical Review Letters*, 114(19):195003, may 2015.

[92] J. P. Couperus, R. Pausch, A. Köhler, O. Zarini, J. M. Krämer, M. Garten, A. Huebl, R. Gebhardt, U. Helbig, S. Bock, K. Zeil, A. Debus, M. Bussmann, U. Schramm, and A. Irman. Demonstration of a beam loaded nanocoulomb-class laser wakefield accelerator. *Nature Communications*, 8(1):487, dec 2017.

[93] Max Tabak, James Hammer, Michael E. Glinsky, William L. Kruer, Scott C. Wilks, John Woodworth, E. Michael Campbell, Michael D. Perry, and Rodney J. Mason. Ignition and high gain with ultrapowerful lasers. *Physics of Plasmas*, 1(5):1626–1634, may 1994.

[94] J. R. Davies. Laser absorption by overdense plasmas in the relativistic regime. *Plasma Physics and Controlled Fusion*, 51(1):014006, jan 2009.

[95] S. C. Wilks, W. L. Kruer, M. Tabak, and A. B. Langdon. Absorption of ultra-intense laser pulses. *Physical Review Letters*, 69(9):1383–1386, aug 1992.

[96] A. Pukhov and J. Meyer-ter Vehn. Laser Hole Boring into Overdense Plasma and Relativistic Electron Currents for Fast Ignition of ICF Targets. *Physical Review Letters*, 79(14):2686–2689, oct 1997.

[97] J. May, J. Tonge, F. Fiuza, R. A. Fonseca, L. O. Silva, C. Ren, and W. B. Mori. Mechanism of generating fast electrons by an intense laser at a steep overdense interface. *Physical Review E*, 84(2):025401(R), aug 2011.

[98] A. J. Kemp and S. C. Wilks. Direct electron acceleration in multi-kilojoule, multi-picosecond laser pulses. *Physics of Plasmas*, 27(10):103106, oct 2020.

[99] S. C. Wilks, A. B. Langdon, T. E. Cowan, M. Roth, M. Singh, S. Hatchett, M. H. Key, D. Pennington, A. MacKinnon, and R. A. Snavely. Energetic proton generation in ultra-intense laser–solid interactions. *Physics of Plasmas*, 8(2):542–549, feb 2001.

[100] P. Mora. Plasma Expansion into a Vacuum. *Physical Review Letters*, 90(18):185002, may 2003.

[101] Y. Sentoku, K. Mima, H. Ruhl, Y. Toyama, R. Kodama, and T. E. Cowan. Laser light and hot electron micro focusing using a conical target. *Physics of Plasmas*, 11(6):3083–3087, jun 2004.

[102] D. R. Rusby, P. M. King, A. Pak, N. Lemos, S. Kerr, G. Cochran, I. Pagano, A. Hannasch, H. Quevedo, M. Spinks, M. Donovan, A. Link, A. Kemp, S. C. Wilks, G. J. Williams, M. J.E. Manuel, Z. Gavin, A. Haid, F. Albert, M. Aufderheide, H. Chen, C. W. Siders, A. Macphee, and A. Mackinnon. Enhancements in laser-generated hot-electron production via focusing cone targets at short pulse and high contrast. *Physical Review E*, 103(5):053207, may 2021.

[103] Y. Ping, R. Shepherd, B. F. Lasinski, M. Tabak, H. Chen, H. K. Chung, K. B. Fournier, S. B. Hansen, A. Kemp, D. A. Liedahl, K. Widmann, S. C. Wilks, W. Rozmus, and M. Sherlock. Absorption of Short Laser Pulses on Solid Targets in the Ultrarelativistic Regime. *Physical Review Letters*, 100(8):085004, feb 2008.

[104] G. J. Williams, A. Link, M. Sherlock, D. A. Alessi, M. Bowers, B. P. Golick, M. Hamamoto, M. R. Hermann, D. Kalantar, K. N. LaFortune, A. J. Mackinnon, A. MacPhee, M. J.E. Manuel, D. Martinez, M. Mauldin, L. Pelz, M. Prantil, M. Quinn, B. Remington, R. Sigurdsson, P. Wegner, K. Youngblood, and Hui Chen. Order-of-magnitude increase in laser-target coupling at near-relativistic intensities using compound parabolic concentrators. *Physical Review E*, 103(3):L031201, mar 2021.

[105] A. Compant La Fontaine. Photon dose produced by a high-intensity laser on a solid target. *Journal of Physics D: Applied Physics*, 47(32):325201, aug 2014.

[106] J Derouillat, A Beck, F Pérez, T Vinci, M Chiaramello, A Grassi, M Flé, G Bouchard, I Plotnikov, N Aunai, J Dargent, C Riconda, and M Grech. Smilei : A collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation. *Computer Physics Communications*, 222:351–373, 2018.

[107] K. J. Bowers, B. J. Albright, L. Yin, B. Bergen, and T. J.T. Kwan. Ultrahigh performance three-dimensional electromagnetic relativistic kinetic plasma simulation. *Physics of Plasmas*, 15(5), 2008.

[108] David P. Grote, Alex Friedman, Jean Luc Vay, and Irving Haber. The WARP code: Modeling high intensity ion beams. In *AIP Conference Proceedings*, volume 749, pages 55–58, 2005.

[109] C. Othmer and J. Schüle. Dynamic load balancing of plasma particle-in-cell simulations: The taskfarm alternative. *Computer Physics Communications*, 2002.

[110] Robert D. Ferraro, Paulett C. Liewer, and Viktor K. Decyk. Dynamic Load Balancing for a 2D Concurrent Plasma PIC Code. *Journal of Computational Physics*, 109(2):329–341, dec 1993.

[111] Igor Surmin, Alexei Bashinov, Sergey Bastrakov, Evgeny Efimenko, Arkady Gonoskov, and Iosif Meyerov. Dynamic load balancing based on rectilinear partitioning in particle-in-cell plasma simulation. In *Lecture Notes in Computer Science (including subseries*

Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), pages 107–119, 2015.

[112] Erik Saule, Erdeniz Ö. Baş, and Ümit V. Çatalyürek. Load-balancing spatially located computations using rectangular partitions. *Journal of Parallel and Distributed Computing*, 72(10):1201–1214, oct 2012.

[113] Kyle G. Miller, Roman P. Lee, Adam Tableman, Anton Helm, Ricardo A. Fonseca, Viktor K. Decyk, and Warren B. Mori. Dynamic load balancing with enhanced shared-memory parallelism for particle-in-cell codes. *Computer Physics Communications*, 259:107633, feb 2021.

[114] Bongki Moon, H.V. Jagadish, Christos Faloutsos, and J.H. Saltz. Analysis of the clustering properties of the Hilbert space-filling curve. *IEEE Transactions on Knowledge and Data Engineering*, 13(1):124–141, 2001.

[115] David Hilbert. Ueber die stetige Abbildung einer Linie auf ein Flähenstück. (On Continuous Mapping of a Line onto a Planar Surface). *Mathematische Annalen*, 38:459–460, 1891.

[116] Xianglong Kong, Michael C. Huang, Chuang Ren, and Viktor K. Decyk. Particle-in-cell simulations with charge-conserving current deposition on graphic processing units. *Journal of Computational Physics*, 230(4):1676–1685, feb 2011.

[117] C. M. Chen, T. Norimatsu, Y. Izawa, T. Yamanaka, and S. Nakai. Heating uniformity of a microwave discharge plasma to redistribute a solid fuel layer inside a cryogenic target for inertial confinement fusion. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 13(6):2908–2913, nov 1995.

[118] R. Betti and O. A. Hurricane. Inertial-confinement fusion with lasers. *Nature Physics*, 12(5):435–448, may 2016.

[119] 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 Physics and Controlled Fusion*, 60(3):034001, mar 2018.

[120] H. Suk, N. Barov, J. B. Rosenzweig, and E. Esarey. Plasma electron trapping and acceleration in a plasma wake field using a density transition. *Physical Review Letters*, 86(6):1011–1014, feb 2001.

[121] E. Oz, S. Deng, T. Katsouleas, P. Muggli, C. D. Barnes, I. Blumenfeld, F. J. Decker, P. Emma, M. J. Hogan, R. Ischebeck, R. H. Iverson, N. Kirby, P. Krejcik, C. O'connell, R. H. Siemann, D. Walz, D. Auerbach, C. E. Clayton, C. Huang, D. K. Johnson, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, and M. Zhou. Ionization-induced electron

trapping in ultrarelativistic plasma wakes. *Physical Review Letters*, 98(8):084801, feb 2007.

[122] A. Pak, K. A. Marsh, S. F. Martins, W. Lu, W. B. Mori, and C. Joshi. Injection and Trapping of Tunnel-Ionized Electrons into Laser-Produced Wakes. *Physical Review Letters*, 104(2):025003, jan 2010.

[123] X. L. Xu, F. Li, W. An, T. N. Dalichaouch, P. Yu, W. Lu, C. Joshi, and W. B. Mori. High quality electron bunch generation using a longitudinal density-tailored plasma-based accelerator in the three-dimensional blowout regime. *Physical Review Accelerators and Beams*, 20(11):111303, nov 2017.

[124] T. N. Dalichaouch, X. L. Xu, F. Li, A. Tableman, F. S. Tsung, W. An, and W. B. Mori. Generating high quality ultrarelativistic electron beams using an evolving electron beam driver. *Physical Review Accelerators and Beams*, 23(2):021304, feb 2020.

[125] A.F. Lifschitz, X. Davoine, E. Lefebvre, J. Faure, C. Rechatin, and V. Malka. Particle-in-Cell modelling of laser–plasma interaction using Fourier decomposition. *Journal of Computational Physics*, 228(5):1803–1814, mar 2009.

[126] Viktor K Decyk. Skeleton Particle-in-Cell Codes on Emerging Computer Architectures. *Computing in Science & Engineering*, 17(2):47–52, mar 2015.

[127] Rémi Lehe, Manuel Kirchen, Igor A. Andriyash, Brendan B. Godfrey, and Jean-Luc Vay. A spectral, quasi-cylindrical and dispersion-free Particle-In-Cell algorithm. *Computer Physics Communications*, 203:66–82, jun 2016.

[128] Gerrit Mur. Absorbing Boundary Conditions for the Finite-Difference Approximation of the Time-Domain Electromagnetic-Field Equations. *IEEE Transactions on Electromagnetic Compatibility*, EMC-23(4):377–382, nov 1981.

[129] Jan De Moerloose and M.A. Stuchly. An efficient way to compare ABCs. *IEEE Antennas and Propagation Magazine*, 38(1):71–75, 1996.

[130] Alvin Bayliss and Eli Turkel. Radiation boundary conditions for wave-like equations. *Communications on Pure and Applied Mathematics*, 33(6):707–725, nov 1980.

[131] Allen Taflove and Susan C Hagness. *Computational Electrodynamics: The Finite-Difference Time-Domain Method*. Artech House, Norwood, MA, 3rd edition, 2005.

[132] C. D. Decker, W. B. Mori, and T. Katsouleas. Particle-in-cell simulations of Raman forward scattering from short-pulse high-intensity lasers. *Physical Review E*, 50(5):R3338–R3341, nov 1994.

[133] F Albert, N Lemos, J. L. Shaw, B. B. Pollock, C Goyon, W Schumaker, A. M. Saunders, K. A. Marsh, A Pak, J. E. Ralph, J. L. Martins, L. D. Amorim, R. W. Falcone, S. H. Glenzer, J. D. Moody, and C Joshi. Observation of Betatron X-Ray Radiation in a Self-Modulated Laser Wakefield Accelerator Driven with Picosecond Laser Pulses. *Physical Review Letters*, 118(13):134801, mar 2017.

[134] Tianhong Wang, Vladimir Khudik, Alexey Arefiev, and Gennady Shvets. Direct laser acceleration of electrons in the plasma bubble by tightly focused laser pulses. *Physics of Plasmas*, 26(8):083101, aug 2019.

[135] A. E. Hussein, A. V. Arefiev, T. Batson, H. Chen, R. S. Craxton, A. S. Davies, D. H. Froula, Z. Gong, D. Haberberger, Y. Ma, P. M. Nilson, W. Theobald, T. Wang, K. Weichman, G. J. Williams, and L. Willingale. Towards the optimisation of direct laser acceleration. *New Journal of Physics*, 23(2):023031, feb 2021.

[136] J L Shaw, F S Tsung, N Vafaei-Najafabadi, K A Marsh, N Lemos, W B Mori, and C Joshi. Role of direct laser acceleration in energy gained by electrons in a laser wakefield accelerator with ionization injection. *Plasma Physics and Controlled Fusion*, 56(8):084006, aug 2014.

[137] A. V. Arefiev, A. P. L. Robinson, and V. N. Khudik. Novel aspects of direct laser acceleration of relativistic electrons. *Journal of Plasma Physics*, 81(4):475810404, aug 2015.

[138] Vladimir Khudik, Alexey Arefiev, Xi Zhang, and Gennady Shvets. Universal scalings for laser acceleration of electrons in ion channels. *Physics of Plasmas*, 23(10):103108, oct 2016.

[139] Vladimir N. Khudik, Xi Zhang, Tianhong Wang, and Gennady Shvets. Far-field constant-gradient laser accelerator of electrons in an ion channel. *Physics of Plasmas*, 25(8):083101, aug 2018.

[140] J L Shaw, N Lemos, K A Marsh, D H Froula, and C Joshi. Experimental signatures of direct-laser-acceleration-assisted laser wakefield acceleration. *Plasma Physics and Controlled Fusion*, 60(4):44012, 2018.

[141] Patrick Mora and Thomas M. Antonsen, Jr. Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas. *Physics of Plasmas*, 4(1):217–229, jan 1997.

[142] Kanani K. M. Lee, L. Robin Benedetti, Raymond Jeanloz, Peter M. Celliers, Jon H. Eggert, Damien G. Hicks, Stephen J. Moon, Andrew Mackinnon, Luis B. Da Silva, David K. Bradley, Walter Unites, Gilbert W. Collins, Emeric Henry, Michel Koenig, Alessandra Benuzzi-Mounaix, John Pasley, and David Neely. Laser-driven shock experiments on precompressed water: Implications for "icy" giant planets. *The Journal of Chemical Physics*, 125(1):014701, jul 2006.

[143] S. V. Bulanov, T. Zh Esirkepov, D. Habs, F. Pegoraro, and T. Tajima. Relativistic laser-matter interaction and relativistic laboratory astrophysics. *The European Physical Journal D*, 55(2):483–507, nov 2009.

[144] N Lemos, F Albert, J L Shaw, D Papp, R Polanek, P King, A L Milder, K A Marsh, A Pak, B B Pollock, B M Hegelich, J D Moody, J Park, R Tommasini, G J Williams, Chen Hui, and C Joshi. Bremsstrahlung hard x-ray source driven by an electron beam from a self-modulated laser wakefield accelerator. *Plasma Physics and Controlled Fusion*, 60(5):54008, 2018.

[145] P. M. King, N. Lemos, J. L. Shaw, A. L. Milder, K. A. Marsh, A. Pak, B. M. Hegelich, P. Michel, J. Moody, C. Joshi, and F. Albert. X-ray analysis methods for sources from self-modulated laser wakefield acceleration driven by picosecond lasers. *Review of Scientific Instruments*, 90(3):033503, mar 2019.

[146] D Gordon, K C Tzeng, C E Clayton, A E Dangor, V Malka, K A Marsh, A Modena, W B Mori, P Muggli, Z Najmudin, D Neely, C Danson, and C Joshi. Observation of Electron Energies Beyond the Linear Dephasing Limit from a Laser-Excited Relativistic Plasma Wave. *Physical Review Letters*, 80(10):2133–2136, 1998.

[147] V. Malka, S. Fritzler, E. Lefebvre, M.-M. Aleonard, F. Burgy, J.-P. Chambaret, J.-F. Chemin, K. Krushelnick, G. Malka, S. P. D. Mangles, Z. Najmudin, M. Pittman, J.-P. Rousseau, J.-N. Scheurer, B. Walton, and A. E. Dangor. Electron Acceleration by a Wake Field Forced by an Intense Ultrashort Laser Pulse. *Science*, 298(5598):1596–1600, nov 2002.

[148] Masahiro Adachi, Eisuke Miura, Susumu Kato, Kazuyoshi Koyama, Shin-ichi Masuda, Takayuki Watanabe, Hiromi Okamoto, Atsushi Ogata, and Mitsumori Tanimoto. Cascade Acceleration of Electrons by Laser Wakefield and Direct Laser Field. *Japanese Journal of Applied Physics*, 45(5A):4214–4218, may 2006.

[149] C. E. Clayton, J. E. Ralph, F. Albert, R. A. Fonseca, S. H. Glenzer, C. Joshi, W. Lu, K. A. Marsh, S. F. Martins, W. B. Mori, A. Pak, F. S. Tsung, B. B. Pollock, J. S. Ross, L. O. Silva, and D. H. Froula. Self-Guided Laser Wakefield Acceleration beyond 1 GeV Using Ionization-Induced Injection. *Physical Review Letters*, 105(10):105003, sep 2010.

[150] L. Willingale, A. V. Arefiev, G. J. Williams, H. Chen, F. Dollar, A. U. Hazi, A. Maksimchuk, M. J.E. Manuel, E. Marley, W. Nazarov, T. Z. Zhao, and C. Zulick. The unexpected role of evolving longitudinal electric fields in generating energetic electrons in relativistically transparent plasmas. *New Journal of Physics*, 20(9):093024, sep 2018.

[151] Guo-Zheng Sun, Edward Ott, Y. C. Lee, and Parvez Guzdar. Self-focusing of short intense pulses in plasmas. *Physics of Fluids*, 30(2):526, 1987.

[152] S. F. Martins, R. A. Fonseca, W. Lu, W. B. Mori, and L. O. Silva. Exploring laser-wakefield-accelerator regimes for near-term lasers using particle-in-cell simulation in Lorentz-boosted frames. *Nature Physics*, 6(4):311–316, apr 2010.

[153] A. V. Arefiev, V. N. Khudik, A. P. L. Robinson, G. Shvets, L. Willingale, and M. Schollmeier. Beyond the ponderomotive limit: Direct laser acceleration of relativistic electrons in sub-critical plasmas. *Physics of Plasmas*, 23(5):056704, may 2016.

[154] J.-L. Vay, R. Lehe, H. Vincenti, B.B. Godfrey, I. Haber, and P. Lee. Recent advances in high-performance modeling of plasma-based acceleration using the full PIC method. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 829:353–357, sep 2016.

[155] Fei Li, Peicheng Yu, Xinlu Xu, Frederico Fiuza, Viktor K Decyk, Thamine Dalichaouch, Asher Davidson, Adam Tableman, Weiming An, Frank S Tsung, Ricardo A Fonseca, Wei Lu, and Warren B Mori. Controlling the numerical Cerenkov instability in PIC simulations using a customized finite difference Maxwell solver and a local FFT based current correction. *Computer Physics Communications*, 214:6–17, 2017.

[156] Károly Németh, Baifei Shen, Yuelin Li, Hairong Shang, Robert Crowell, Katherine C. Harkay, and John R. Cary. Laser-Driven Coherent Betatron Oscillation in a Laser-Wakefield Cavity. *Physical Review Letters*, 100(9):095002, mar 2008.

[157] J. T. Mendonça and F. Doveil. Stochasticity in plasmas with electromagnetic waves. *Journal of Plasma Physics*, 28(3):485–493, dec 1982.

[158] Z.-M. Sheng, K. Mima, Y. Sentoku, M. S. Jovanović, T. Taguchi, J. Zhang, and J. Meyer-ter Vehn. Stochastic Heating and Acceleration of Electrons in Colliding Laser Fields in Plasma. *Physical Review Letters*, 88(5):055004, jan 2002.

[159] Zheng-Ming Sheng, Kunioki Mima, Jie Zhang, and Jürgen Meyer-ter Vehn. Efficient acceleration of electrons with counterpropagating intense laser pulses in vacuum and underdense plasma. *Physical Review E*, 69(1):016407, jan 2004.

[160] M Borland. ELEGANT: A flexible SDDS-compliant code for accelerator simulation. Technical Report September 2000, Argonne National Laboratory (ANL), Argonne, IL (United States), aug 2000.

[161] Yujian Zhao, Weiming An, Xinlu Xu, Fei Li, Lance Hildebrand, Mark J. Hogan, Vitaly Yakimenko, Chan Joshi, and Warren B. Mori. Emittance preservation through density ramp matching sections in a plasma wakefield accelerator. *Physical Review Accelerators and Beams*, 23(1):011302, jan 2020.

[162] I. Watts, M. Zepf, E. L. Clark, M. Tatarakis, K. Krushelnick, A. E. Dangor, R. M. Allott, R. J. Clarke, D. Neely, and P. A. Norreys. Dynamics of the Critical Surface in High-Intensity Laser-Solid Interactions: Modulation of the XUV Harmonic Spectra. *Physical Review Letters*, 88(15):155001, mar 2002.

[163] George D. Tsakiris, Klaus Eidmann, Jürgen Meyer-ter Vehn, and Ferenc Krausz. Route to intense single attosecond pulses. *New Journal of Physics*, 8, 2006.

[164] H.-S. Park, D. M. Chambers, H.-K. Chung, R. J. Clarke, R. Eagleton, E. Giraldez, T. Goldsack, R. Heathcote, N. Izumi, M. H. Key, J. A. King, J. A. Koch, O. L. Landen, A. Nikroo, P. K. Patel, D. F. Price, B. A. Remington, H. F. Robey, R. A. Snavely, D. A. Steinman, R. B. Stephens, C. Stoeckl, M. Storm, M. Tabak, W. Theobald, R. P. J. Town, J. E. Wickersham, and B. B. Zhang. High-energy Kα radiography using high-intensity, short-pulse lasers. *Physics of Plasmas*, 13(5):056309, may 2006.

[165] L. L. Ji, A. Pukhov, E. N. Nerush, I. Yu Kostyukov, B. F. Shen, and K. U. Akli. Energy partition, γ-ray emission, and radiation reaction in the near-quantum electrodynamical regime of laser-plasma interaction. *Physics of Plasmas*, 21(2):023109, feb 2014.

[166] R. Capdessus, M. Lobet, E. D'Humières, and V. T. Tikhonchuk. γ -ray generation enhancement by the charge separation field in laser-target interaction in the radiation dominated regime. *Physics of Plasmas*, 21(12):123120, dec 2014.

[167] A P L Robinson, M. Zepf, S. Kar, R. G. Evans, and C. Bellei. Radiation pressure acceleration of thin foils with circularly polarized laser pulses. *New Journal of Physics*, 10(1):013021, jan 2008.

[168] Anatoly Spitkovsky. On the Structure of Relativistic Collisionless Shocks in Electron-Ion Plasmas. *The Astrophysical Journal*, 673(1):L39–L42, jan 2008.

[169] S. F. Martins, R. A. Fonseca, L. O. Silva, and W. B. Mori. ION DYNAMICS AND ACCELERATION IN RELATIVISTIC SHOCKS. *The Astrophysical Journal*, 695(2):L189–L193, apr 2009.

[170] F. Fiuza, R. A. Fonseca, J. Tonge, W. B. Mori, and L. O. Silva. Weibel-Instability-Mediated Collisionless Shocks in the Laboratory with Ultraintense Lasers. *Physical Review Letters*, 108(23):235004, jun 2012.

[171] J. Tonge, J. May, W. B. Mori, F. Fiuza, S. F. Martins, R. A. Fonseca, L. O. Silva, and C. Ren. A simulation study of fast ignition with ultrahigh intensity lasers. *Physics of Plasmas*, 16(5):056311, may 2009.

[172] G. Malka and J. L. Miquel. Experimental Confirmation of Ponderomotive-Force Electrons Produced by an Ultrarelativistic Laser Pulse on a Solid Target. *Physical Review Letters*, 77(1):75–78, jul 1996.

[173] Scott C. Wilks and William L. Kruer. Absorption of ultrashort, ultra-intense laser light by solids and overdense plasmas. *IEEE Journal of Quantum Electronics*, 33(11):1954–1968, 1997.

[174] A. Pukhov and J. Meyer-ter Vehn. Relativistic laser-plasma interaction by multi-dimensional particle-in-cell simulations. *Physics of Plasmas*, 5(5):1880–1886, may 1998.

[175] M. I. K. Santala, E. Clark, I. Watts, F. N. Beg, M. Tatarakis, M. Zepf, K. Krushelnick, A. E. Dangor, T. McCanny, I. Spencer, R. P. Singhal, K. W. D. Ledingham, S. C. Wilks, A. C. Machacek, J. S. Wark, R. Allott, R. J. Clarke, and P. A. Norreys. Effect of the Plasma Density Scale Length on the Direction of Fast Electrons in Relativistic Laser-Solid Interactions. *Physical Review Letters*, 84(7):1459–1462, feb 2000.

[176] J.C. Kieffer, A. Krol, Z. Jiang, C.C. Chamberlain, E. Scalzetti, and Z. Ichalalene. Future of laser-based X-ray sources for medical imaging. *Applied Physics B*, 74(S1):s75–s81, jun 2002.

[177] O. Culfa, G. J. Tallents, A. K. Rossall, E. Wagenaars, C. P. Ridgers, C. D. Murphy, R. J. Dance, R. J. Gray, P. McKenna, C. D. R. Brown, S. F. James, D. J. Hoarty, N. Booth, A. P. L. Robinson, K. L. Lancaster, S. A. Pikuz, A. Ya Faenov, T. Kampfer, K. S. Schulze, I. Uschmann, and N. C. Woolsey. Plasma scale-length effects on electron energy spectra in high-irradiance laser plasmas. *Physical Review E*, 93(4):043201, apr 2016.

[178] Sheng Jiang, Andrew G. Krygier, Douglass W. Schumacher, Kramer U. Akli, and Richard R. Freeman. Enhancing Bremsstrahlung production from ultraintense laser-solid interactions with front surface structures. *The European Physical Journal D*, 68(10):283, oct 2014.

[179] Tina Ebert, Nico W. Neumann, Leonard N. K. Döhl, Jonathan Jarrett, Christopher Baird, Robert Heathcote, Markus Hesse, Aasia Hughes, Paul McKenna, David Neely, Dean Rusby, Gabriel Schaumann, Christopher Spindloe, Alexandra Tebartz, Nigel Woolsey, and Markus Roth. Enhanced brightness of a laser-driven x-ray and particle source by microstructured surfaces of silicon targets. *Physics of Plasmas*, 27(4):043106, apr 2020.

[180] S. A. Gaillard, T. Kluge, K. A. Flippo, M. Bussmann, B. Gall, T. Lockard, M. Geissel, D. T. Offermann, M. Schollmeier, Y. Sentoku, and T. E. Cowan. Increased laser-accelerated proton energies via direct laser-light-pressure acceleration of electrons in microcone targets. *Physics of Plasmas*, 18(5):056710, may 2011.

[181] T. Kluge, S. A. Gaillard, K. A. Flippo, T. Burris-Mog, W. Enghardt, B. Gall, M. Geissel, A. Helm, S. D. Kraft, T. Lockard, J. Metzkes, D. T. Offermann, M. Schollmeier, U. Schramm, K. Zeil, M. Bussmann, and T. E. Cowan. High proton energies from cone

targets: Electron acceleration mechanisms. *New Journal of Physics*, 14(2):023038, feb 2012.

[182] R. A. Simpson, G. G. Scott, D. Mariscal, D. Rusby, P. M. King, E. Grace, A. Aghedo, I. Pagano, M. Sinclair, C. Armstrong, M. J.E. Manuel, A. Haid, K. Flippo, L. Winslow, M. Gatu-Johnson, J. A. Frenje, D. Neely, S. Kerr, G. J. Williams, S. Andrews, R. Cauble, K. Charron, R. Costa, B. Fischer, S. Maricle, B. Stuart, F. Albert, N. Lemos, A. Mackinnon, A. MacPhee, A. Pak, and T. Ma. Scaling of laser-driven electron and proton acceleration as a function of laser pulse duration, energy, and intensity in the multi-picosecond regime. *Physics of Plasmas*, 28(1):013108, jan 2021.

[183] J. May, J. Tonge, I. Ellis, W. B. Mori, F. Fiuza, R. A. Fonseca, L. O. Silva, and C. Ren. Enhanced stopping of macro-particles in particle-in-cell simulations. *Physics of Plasmas*, 21(5):052703, 2014.

[184] Kyle G. Miller, Joshua May, Frederico Fiuza, and Warren B. Mori. Extended particle absorber for efficient modeling of intense laser–solid interactions. *Physics of Plasmas*, 28(11):112702, nov 2021.

[185] A R Bell, J R Davies, S Guerin, and H Ruhl. Fast-electron transport in high-intensity short-pulse laser - solid experiments. *Plasma Physics and Controlled Fusion*, 39(5):653–659, may 1997.

[186] M. C. Levy, S. C. Wilks, M. Tabak, and M. G. Baring. Conservation laws and conversion efficiency in ultraintense laser-overdense plasma interactions. *Physics of Plasmas*, 20(10):103101, oct 2013.

[187] J. C. Adam, A. Héron, and G. Laval. Dispersion and Transport of Energetic Particles due to the Interaction of Intense Laser Pulses with Overdense Plasmas. *Physical Review Letters*, 97(20):205006, nov 2006.

[188] Joshua May. *On the Acceleration and Transport of Electrons Generated by Intense Laser-Plasma Interactions at Sharp Interfaces*. PhD thesis, University of California, Los Angeles, 2017.

[189] Stephen P. Hatchett, Curtis G. Brown, Thomas E. Cowan, Eugene A. Henry, Joy S. Johnson, Michael H. Key, Jeffrey A. Koch, A. Bruce Langdon, Barbara F. Lasinski, Richard W. Lee, Andrew J. Mackinnon, Deanna M. Pennington, Michael D. Perry, Thomas W. Phillips, Markus Roth, T. Craig Sangster, Mike S. Singh, Richard A. Snavely, Mark A. Stoyer, Scott C. Wilks, and Kazuhito Yasuike. Electron, photon, and ion beams from the relativistic interaction of Petawatt laser pulses with solid targets. *Physics of Plasmas*, 7(5):2076–2082, may 2000.

[190] A. J. Mackinnon, Y. Sentoku, P. K. Patel, D. W. Price, S. Hatchett, M. H. Key, C. Andersen, R. Snavely, and R. R. Freeman. Enhancement of Proton Acceleration

by Hot-Electron Recirculation in Thin Foils Irradiated by Ultraintense Laser Pulses. *Physical Review Letters*, 88(21):215006, may 2002.

[191] Claire Ellen Max, Wallace M. Manheimer, and J. J. Thomson. Enhanced transport across laser generated magnetic fields. *Physics of Fluids*, 21(1):128, 1978.

[192] D. W. Forslund and J. U. Brackbill. Magnetic-Field-Induced Surface Transport on Laser-Irradiated Foils. *Physical Review Letters*, 48(23):1614–1617, jun 1982.

[193] F. S. Tsung, W. Lu, M. Tzoufras, W. B. Mori, C. Joshi, J. M. Vieira, L. O. Silva, and R. A. Fonseca. Simulation of monoenergetic electron generation via laser wakefield accelerators for 5–25TW lasers. *Physics of Plasmas*, 13(5):056708, may 2006.

[194] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytracek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell'Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani, D. Gibin, J.J. Gómez Cadenas, I. González, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F.W. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata, M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampén, V. Lara, V. Lefebure, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. Mora de Freitas, Y. Morita, K. Murakami, M. Nagamatu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O'Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M.G. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. Di Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. Safai Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J.P. Wellisch, T. Wenaus, D.C. Williams, D. Wright, T. Yamada, H. Yoshida, and D. Zschiesche. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, jul 2003.

[195] Scott C. Wilks. Simulations of ultraintense laser–plasma interactions. *Physics of Fluids B: Plasma Physics*, 5(7):2603–2608, jul 1993.

[196] D. W. Forslund, J. M. Kindel, W. B. Mori, C. Joshi, and J. M. Dawson. Two-Dimensional Simulations of Single-Frequency and Beat-Wave Laser-Plasma Heating. *Physical Review Letters*, 54(6):558–561, feb 1985.

[197] J. Meyer-ter Vehn and Z. M. Sheng. On electron acceleration by intense laser pulses in the presence of a stochastic field. *Physics of Plasmas*, 6(3):641–644, mar 1999.

[198] A. Sorokovikova, A. V. Arefiev, C. McGuffey, B. Qiao, A. P. L. Robinson, M. S. Wei, H. S. McLean, and F. N. Beg. Generation of Superponderomotive Electrons in Multipicosecond Interactions of Kilojoule Laser Beams with Solid-Density Plasmas. *Physical Review Letters*, 116(15):155001, apr 2016.

[199] Barbara F. Lasinski, A. Bruce Langdon, Stephen P. Hatchett, Michael H. Key, and Max Tabak. Particle-in-cell simulations of ultra intense laser pulses propagating through overdense plasma for fast-ignitor and radiography applications. *Physics of Plasmas*, 6(5):2041–2047, may 1999.

[200] L. G. Huang, H. Takabe, and T. E. Cowan. Maximizing magnetic field generation in high power laser–solid interactions. *High Power Laser Science and Engineering*, 7:e22, apr 2019.

[201] N. Shukla, K. Schoeffler, E. Boella, J. Vieira, R. Fonseca, and L. O. Silva. Interplay between the Weibel instability and the Biermann battery in realistic laser-solid interactions. *Physical Review Research*, 2(2):023129, may 2020.

[202] A. Bret, L. Gremillet, and D. Bénisti. Exact relativistic kinetic theory of the full unstable spectrum of an electron-beam–plasma system with Maxwell-Jüttner distribution functions. *Physical Review E*, 81(3):036402, mar 2010.

[203] J. L. Shaw, M. A. Romo-Gonzalez, N. Lemos, P. M. King, G. Bruhaug, K. G. Miller, C Dorrer, B Kruschwitz, L Waxer, G J Williams, M V Ambat, M. M. McKie, M. D. Sinclair, W. B. Mori, C. Joshi, Hui Chen, J. P. Palastro, F. Albert, and D. H. Froula. Microcoulomb ($0.7 \pm \frac{0.4}{0.2}$ μC) laser plasma accelerator on OMEGA EP. *Scientific Reports*, 11(1):7498, dec 2021.

[204] Luís O. Silva, Michael Marti, Jonathan R. Davies, Ricardo A. Fonseca, Chuang Ren, Frank S. Tsung, and Warren B. Mori. Proton Shock Acceleration in Laser-Plasma Interactions. *Physical Review Letters*, 92(1):015002, jan 2004.

[205] Y. Sentoku, K. Mima, P. Kaw, and K. Nishikawa. Anomalous Resistivity Resulting from MeV-Electron Transport in Overdense Plasma. *Physical Review Letters*, 90(15):155001, apr 2003.

[206] A. Pukhov. Three-Dimensional Simulations of Ion Acceleration from a Foil Irradiated by a Short-Pulse Laser. *Physical Review Letters*, 86(16):3562–3565, apr 2001.

[207] A. J. Kemp and L. Divol. Interaction Physics of Multipicosecond Petawatt Laser Pulses with Overdense Plasma. *Physical Review Letters*, 109(19):195005, nov 2012.

[208] T.-Y. Brian Yang, William L. Kruer, Richard M. More, and A. Bruce Langdon. Absorption of laser light in overdense plasmas by sheath inverse bremsstrahlung. *Physics of Plasmas*, 2(8):3146–3154, aug 1995.

[209] A.J. Kemp, F. Fiuza, A. Debayle, T. Johzaki, W.B. Mori, P.K. Patel, Y. Sentoku, and L.O. Silva. Laser–plasma interactions for fast ignition. *Nuclear Fusion*, 54(5):054002, may 2014.

[210] Y. Sentoku, T. E. Cowan, A. Kemp, and H. Ruhl. High energy proton acceleration in interaction of short laser pulse with dense plasma target. *Physics of Plasmas*, 10(5):2009–2015, may 2003.

[211] M. N. Quinn, X. H. Yuan, X. X. Lin, D. C. Carroll, O. Tresca, R. J. Gray, M. Coury, C Li, Y. T. Li, C. M. Brenner, A P L Robinson, D. Neely, B. Zielbauer, B. Aurand, J. Fils, T. Kuehl, and P. McKenna. Refluxing of fast electrons in solid targets irradiated by intense, picosecond laser pulses. *Plasma Physics and Controlled Fusion*, 53(2):025007, feb 2011.

[212] P. M. Nilson, W. Theobald, J. F. Myatt, C. Stoeckl, M. Storm, J. D. Zuegel, R. Betti, D. D. Meyerhofer, and T. C. Sangster. Bulk heating of solid-density plasmas during high-intensity-laser plasma interactions. *Physical Review E*, 79(1):016406, jan 2009.

[213] C W Allen. *Astrophysical quantities*. University of London, Athlone Press, London, UK, 3rd edition, 1973.

[214] Merran Evans, Nicholas Hastings, and Brian Peacock. Terms and Symbols. In *Statistical Distributions*, chapter 2, pages 3–16. Wiley-Interscience, New York, NY, 3rd edition, 2000.