

UCLA

UCLA Electronic Theses and Dissertations

Title

Wakes in Inertial Fusion Plasmas

Permalink

<https://escholarship.org/uc/item/5wb109v8>

Author

Ellis, Ian Norman

Publication Date

2014

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Wakes in Inertial Fusion Plasmas

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

by

Ian Norman Ellis

2014

© Copyright by
Ian Norman Ellis
2014

ABSTRACT OF THE DISSERTATION

Wakes in Inertial Fusion Plasmas

by

Ian Norman Ellis

Doctor of Philosophy in Physics

University of California, Los Angeles, 2014

Professor Warren B. Mori, Chair

Plasma wave wakes, which are the collective oscillatory response near the plasma frequency to the propagation of particles or electromagnetic waves through a plasma, play a critical role in many plasma processes. New results from backwards stimulated Raman scattering (BSRS), in which wakes with phase velocities much less than the speed of light are induced by the beating of counter-propagating light waves, and from electron beam stopping, in which the wakes are produced by the motion of relativistically propagating electrons through the dense plasma, are discussed. Both processes play important roles in Inertial Confinement Fusion (ICF). In BSRS, laser light is scattered backwards out of the plasma, decreasing the energy available to compress the ICF capsule and affecting the symmetry of where the laser energy hits the hohlraum wall in indirect drive ICF. The plasma wave wake can also generate superthermal electrons that can preheat the core and/or the ablator. Electron beam stopping plays a critical role in the Fast Ignition (FI) ICF concept, in which a beam of relativistic electrons is used to heat the target core to ignition temperatures after the compression stage. The beam stopping power determines the effectiveness of the heating process. This dissertation covers new discoveries on the importance of plasma wave wakes in both BSRS and electron beam stopping.

In the SRS studies, 1D particle-in-cell (PIC) simulations using OSIRIS are per-

formed, which model a short-duration ($\sim 500\omega_0^{-1}$ FWHM) counter-propagating scattered light seed pulse in the presence of a constant pump laser with an intensity far below the absolute instability threshold for plasma waves undergoing Landau damping. The seed undergoes linear convective Raman amplification and dominates over the amplification of fluctuations due to particle discreteness. The simulation results are in good agreement with results from a coupled-mode solver when special relativity and the effects of finite size PIC simulation particles are accounted for. Linear gain spectra including both effects are discussed. Extending the PIC simulations past when the seed exits the simulation domain reveals bursts of large-amplitude scattering in many cases, which do not occur in simulations without the seed pulse. These bursts can have amplitudes several times greater than the amplified seed pulse, and an examination of the orbits of particles trapped in the wake illustrates that the bursts are caused by a reduction of Landau damping due to particle trapping. This large-amplitude scattering is caused by the seed inducing a wake earlier in the simulation, thus modifying the distribution function. Performing simulations with longer duration seeds leads to parts of the seeds reaching amplitudes several times more than the steady-state linear theory results, similarly caused by a reduction of Landau damping. Simulations with continuous seeds demonstrate that the onset of inflation depends on the seed wavelength and incident intensity, and oscillations in the reflectivity are observed at a frequency equal to the difference between the seed frequency and the frequency at which the inflationary SRS grows.

In the electron beam stopping studies, 3D PIC simulations are performed of relativistic electrons with a momentum of $10m_e c$ propagating in a cold FI core plasma. Some of the simulations use one simulation particle per real particle, and particle sizes much smaller than the interparticle spacing. The wake made by a single electron is compared against that calculated using cold fluid theory assuming the phase velocity of the wake is near the speed of light. The results agree for

the first wavelength of the wake. However, the shape of the wake changes for succeeding wavelengths and depends on the background plasma temperature, with the concavity pointing in the direction the electron is moving in cold plasmas and in the opposite direction as the plasma temperature increases. In the warm plasma the curvature is described by electrostatic Vlasov theory (for $v_{particle} \gg v_{th}$) and is due to the diffraction of the wave, while for cold plasmas the curvature is due to nonlinear radial oscillations of background electrons. Beams with multiple electrons exhibit correlation effects caused by electrons interacting through their wakes. Non-divergent beams are simulated, and a significant time-dependent increase in the stopping power is observed when the average electron spacing is $2c/\omega_{pe}$ or less. This increase is caused by beam-plasma-like instabilities including self-focusing and/or filamentation and the beam-plasma-like instability. The stopping power growth rate and peak value depend on the beam size and density. For long beams with dimensions of $10c/\omega_{pe} \times 10c/\omega_{pe} \times 80c/\omega_{pe}$ and an inter-particle separation of $0.25c/\omega_{pe}$ ($n_b/n_0 \approx 4 \times 10^{-3}$), the peak stopping power averaged over the electrons is $(1 \pm 3) \times 10^3$ times that of an uncorrelated electron. These results indicate that an enhanced energy-independent or weakly dependent correlated stopping may occur for Fast Ignition scenarios, even for interparticle spacings when discreteness effects are important. The dependence of correlation effects on beam electron separation in terms of c/ω_{pe} also indicates that Fast Ignition may be possible with core densities below those designed using single-electron stopping powers. Target optimization to exploit correlated stopping in the target core may be possible once the effects of angular spread and energy spread are understood. Furthermore, this work begins to allow a connection from the discrete wakes effect to collective instabilities [MF02] as the interparticle spacing is decreased relative to the size of the wake due to the use of denser beams, lower plasma densities, and the filamentation/self-focusing of the beam.

The dissertation of Ian Norman Ellis is approved.

Frank R. Graziani

Viktor K. Decyk

Christoph Niemann

Chandrashekhar J. Joshi

George J. Morales

Warren B. Mori, Committee Chair

University of California, Los Angeles

2014

*To my parents . . .
who have supported me through all
my endeavors in life, and were there
to help me stand back up whenever I
stumbled along the way.*

TABLE OF CONTENTS

1	Introduction	1
2	Convective Raman Amplification of Light Pulses Causing Kinetic Inflation in Inertial Fusion Plasmas	7
2.1	Introduction	7
2.2	Geometry and Plasma Conditions	10
2.3	Linear Theory of Convective BSRS Gain	12
2.3.1	Non-Relativistic Theory	12
2.3.2	Relativistic Modification	14
2.3.3	PIC Modification	16
2.4	Subtraction Technique	20
2.5	Convective Gain Simulations & Coupled-Mode Results	22
2.5.1	OSIRIS Convective Gain Simulations	22
2.5.2	The Coupled-Mode Equations	26
2.6	Inflation After Seed Passage	30
2.7	Inflation of the Seed	42
2.8	Inflation of Continuous Seeds	47
2.9	Conclusions and Future Research	53
3	Studies of Particle Wake Potentials in Plasmas	56
3.1	Introduction	56
3.2	PIC & PPPM Methods	58
3.3	Collisionless Kinetic Theory of Particle Wakes	60

3.4	The Subtraction Technique	61
3.5	PIC Results and Comparison with Collisionless Kinetic Theory . .	63
3.6	PPPM Results and Comparison with Traditional PIC	68
3.7	Fine-Mesh PIC	73
3.8	Conclusions	75
4	Single-Particle Stopping Power	77
4.1	Introduction	77
4.2	Stopping Power Concepts	77
4.3	Simple Relativistic Cold Fluid Wake Stopping Power	81
4.4	Classical Bohr Stopping	85
4.4.1	Peter & Meyer-ter-Vehn Derivation	86
4.4.2	Typical (Collisional) Derivation	90
4.5	Wake and Collision Contributions	92
4.6	The QED Stopping Formula	93
4.7	Comparison with Simulations	96
4.7.1	Comparisons with ddcMD and BEPS	97
4.7.2	Comparison with OSIRIS	109
4.7.3	Comparison with QuickPIC	112
4.8	Conclusion	116
5	The Anatomy of Relativistic Wakes	118
5.1	Introduction	118
5.2	Longitudinal and Focusing Fields	118
5.3	Curvature and Spreading	122

6	Correlated Electron Stopping in Fast Ignition Plasmas	130
6.1	Introduction	130
6.2	Simulation Codes, Parameters, and Simplifying Assumptions . . .	133
6.3	Betatron Motion	135
6.4	Beam Scaling Parameter	139
6.5	Relativistic Fluid Streaming Instabilities in 1D	141
6.6	Transverse Beam Self-Modulation	147
6.7	A Beam of 125 Electrons with Various Densities	149
6.7.1	One Beam Electron Per Cubic Skin Depth	149
6.7.2	Various Beam Electron Densities	155
6.8	Various Beam Densities Initially Inside a 1,000 $(c/\omega_{pe})^3$ Cube . .	157
6.9	Box-Length Beams	165
6.10	Conclusion and Future Work	173
7	Conclusions and Future Work	176
A	QuickPIC	180
A.1	Introduction	180
A.2	Basic Setup and Implementation Details	180
A.3	The Quasi-Static Approximation	181
A.4	The Quantity ψ	183
A.5	Equations of Motion	184
A.6	Field Equations	187
A.7	Conclusion	189
	References	191

LIST OF FIGURES

2.1	The geometry of the OSIRIS simulations.	10
2.2	Several calculations of the linear convective BSRS gain spectra for a box $1790c/\omega_0$ ($100 \mu\text{m}$) long. The dash-dotted curves are analytic results, and the (dashed, solid) ones take into account PIC effects with the (5,2)-pass filter.	19
2.3	The longitudinal electric field in a simulation with (a) and without (b) a seed pulse. Subtracting the latter from the former reveals the plasma wave (c,d).	21
2.4	The shapes of the seed pulses in the simulations. A Gaussian-like pulse (a) with a $500\omega_0^{-1}$ rise and fall time ($\sim 500\omega_0^{-1}$ FWHM), and a flat-top pulse (b) with a $200\omega_0^{-1}$ Gaussian-like rise and fall, and a steady amplitude for $600\omega_0^{-1}$ in between.	23
2.5	Scattered light vs. position and time (a), lineouts of the scattered light vs. position at various times (b), and plasma wave amplitude (c) from a simulation using a Gaussian-like seed pulse with $I_{1s} = 5 \times 10^{-4}I_0$ and $\lambda_1 = 1.644\lambda_0$	24

2.6	Measured gain in simulations with box length $1790c/\omega_0$ ($100 \mu\text{m}$) vs. seed wavelength (a) and seed intensity I_{1s} (b). Plotted are the gains we measure using a Gaussian-like pulse (green with square markers) and a flat-top pulse (blue with circle markers), both with $\lambda_1 = 1.644\lambda_0$. (a) includes the gain from simulations with a Gaussian-like pulse and the 5-pass filter (black with ‘x’ markers), plus those without relativistic effects (brown with ‘+’ markers). Several theoretical gain curves taking into account PIC effects are included: 2-pass non-relativistic is dash-dotted red, 2-pass relativistic is solid magenta, and 5-pass relativistic is dashed purple.	25
2.7	The BSRS reflected light using a flat-top seed with $\lambda_1 = 1.644\lambda_0$ (a) and $1.650\lambda_0$ (b). The simulation results are plotted with solid magenta lines while the coupled-mode results are plotted with dash-dotted black lines. The horizontal dashed red line indicates the maximum seed amplitude with no gain.	28
2.8	The measured gain in a box of length $1790c/\omega_0$ ($100 \mu\text{m}$) predicted by the coupled-mode solver as we vary the wavelength using a Gaussian-like pulse (dashed brown with ‘+’ markers) and a flat-top pulse (dashed magenta with circle markers). For comparison, we include the PIC simulation result for the Gaussian-like pulse (solid blue with ‘x’ markers) and the flat-top pulse (solid red with square markers) along with the linear relativistic gain curve taking into account PIC effects (dash-dotted green).	29
2.9	The longitudinal field vs. space at $t = 4,000\omega_0^{-1}$ when using a flat-top seed with $\lambda_1 = 1.644\lambda_0$ (a) and $1.650\lambda_0$ (b); the same runs as used in Figure 2.7. The simulation results are plotted with solid magenta lines while the coupled-mode results are plotted with dash-dotted black lines.	30

2.10	The scattered light (a) and the plasma wave (b) seen when we extend the duration of the simulation using a Gaussian-like seed pulse with $I_{1s} = 8 \times 10^{-3}I_0$ and $\lambda_1 = 1.644\lambda_0$. The seed exits the box around $t = 5,000\omega_0^{-1}$	32
2.11	Evolution of the scattered light (a) and the longitudinal field (the PIC simulation field is filtered) (b) with time at $x = 550c/\omega_0$ for the simulation of Figure 2.10. The simulation results are plotted with solid blue lines and the coupled-mode results are plotted with dashed black lines.	34
2.12	Orbits of particles trapped in the plasma wave in the simulation of Figure 2.10. The orbits are plotted for $2,000\omega_0^{-1}$ starting at $t_0 = 5,000\omega_0^{-1}$ in the wave frame, where $v_{ph} = 0.271c$ is the plasma wave phase speed and the orbits are centered around the relativistic $p_{ph} = 0.282m_e c$	35
2.13	The distribution functions from the simulation of Figure 2.10 at $t = 6,000\omega_0^{-1}$ (solid blue) and $t = 17,000\omega_0^{-1}$ (dashed black) along with the Jüttner distribution for a 2.5keV electron plasma (dash-dotted red). The measured distributions have a plateau beginning around $\gamma\beta = 0.26$, indicating particle trapping.	36
2.14	The reflected light in extended duration simulations using a Gaussian-like seed pulse with $\lambda_1 = 1.644\lambda_0$ and various maximum initial intensities. Shown are $I_{1s} = 4 \times 10^{-3}I_0$ (solid blue), $8 \times 10^{-3}I_0$ (dotted green), $0.016I_0$ (dash-dotted red), and $0.032I_0$ (dashed black). . .	38

2.15 The scattered light (a) and plasma wave (b) in a simulation using a flat-top seed pulse with $I_{1s} = 8 \times 10^{-3} I_0$ and $\lambda_1 = 1.638\lambda_0$. Below them are the Wigner transforms of the reflected light at $x = 0$ (c), the plasma wave at $x = 500c/\omega_0$ (d), and the plasma wave at $x = 900c/\omega_0$ (e). 40

2.16 The scattered light (a) and plasma wave (b) in a simulation using a flat-top seed pulse with $I_{1s} = 8 \times 10^{-3} I_0$ and $\lambda_1 = 1.627\lambda_0$. Below them are the Wigner transforms of the reflected light at $x = 0$ (c) and the plasma wave at $x = 900c/\omega_0$ (d). 41

2.17 Results from simulations with flat-top seed pulses with $I_{1s} = 8 \times 10^{-3} I_0$ and $\lambda_1 = 1.644\lambda_0$. The scattered light (left) and plasma wave (right) for simulations using a pulse duration of $(1000, 3000)\omega_0^{-1}$ are in panels (a-b, c-d), respectively. 43

2.18 The reflected light seen in several simulations using a flat-top seed of duration $6,000\omega_0^{-1}$ with $I_{1s} = 8 \times 10^{-3} I_0$ for various wavelengths. The two top curves are for seeds with $\lambda_1 = 1.638\lambda_0$ (solid red) and $1.644\lambda_0$ (dashed cyan). The four lower curves are for seeds with $\lambda_1 = 1.627\lambda_0$ (dotted blue), $1.632\lambda_0$ (dashed green), $1.650\lambda_0$ (dash-dotted purple), and $1.658\lambda_0$ (solid yellow). For comparison, we mark the steady-state linear relativistic PIC values using horizontal dashes on the left side of the plot. The red vertical dash on the lower right side of the plot indicates approximately when the seeds end. 44

2.19 The scattered light (a), and the resonance product R_p from Equation 2.34 (b) for the flat-top seed of duration $6,000\omega_0^{-1}$ with $\lambda_1 = 1.644\lambda_0$ 46

2.20	The scattered light (a), and the resonance product R_p from Equation 2.34 (b) for the flat-top seed of duration $6,000\omega_0^{-1}$ with $\lambda_1 = 1.638\lambda_0$	48
2.21	The time-average measured gain seen in simulations as we vary the continuous seed intensity using seeds with $\lambda_1 = 1.644\lambda_0$ (dashed blue) and $1.658\lambda_0$ (solid red). For comparison, we use horizontal dashes on the left side of the plot to mark the steady-state gain from linear relativistic PIC theory.	49
2.22	The reflected light in simulations using a continuous seed with $\lambda_1 = 1.658\lambda_0$ (solid red in Figure 2.21). $I_{1s}/I_0 = 4 \times 10^{-3}$ (a), 8×10^{-3} (b), and $1.024I_0$ (c). The horizontal dash-dotted line indicates the amplitude of the seed (unamplified) reflected light.	52
2.23	The reflected light in simulations using a continuous seed with $\lambda_1 = 1.644\lambda_0$ (dashed blue in Figure 2.21). $I_{1s}/I_0 = 1.25 \times 10^{-4}$ (a), 8×10^{-3} (b), and 1.024 (c). The horizontal dash-dotted line indicates the amplitude of the seed (unamplified) reflected light.	53
3.1	The cycles used in the PIC the PPPM methods. Conceptually, the two cycles are identical except for accounting for pair-wise interactions (red box) in the PPPM method.	59
3.2	The electrostatic potential (units e/λ_{De}) from the BEPS simulation of Section 3.5, with a test charge and an additional neutralizing charge (left), one without the test charge and its neutralizing charge but otherwise identical (center), and the difference of the two (right). The test charge is at $\vec{x} = (128\lambda_{De}, 128\lambda_{De}, 214\lambda_{De})$ at the time shown.	62

3.3 The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in the BEPS simulation of Section 3.5. The z position of the test charge is $91\lambda_{De}$ and $334\lambda_{De}$ in the left and right plots, respectively. 65

3.4 The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in the BEPS simulation of Section 3.5 with various test particle speeds. The test particle speed is $3v_{th}$ and $10v_{th}$ and its z position is $154\lambda_{De}$ and $364\lambda_{De}$ in the left and right plots, respectively. 66

3.5 The electrostatic potential along the trajectory of the test charge with speed $3v_{th}$ in the BEPS simulation of Section 3.5 (solid) and as predicted by collisionless kinetic theory (dotted) at $20\omega_{pe}^{-1}$. The vertical line indicates the position of the test charge. 67

3.6 The electrostatic potential along the trajectory of the test particle with speed $3v_{th}$ at various times in the BEPS simulation of Section 3.5 (solid) and as predicted by collisionless kinetic theory taking into account finite particle size and Fourier space smoothing (dash-dotted). The vertical lines indicate the position of the test charge. 69

3.7 The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in a ddcMD simulation. The z position of the test charge is $0.105\lambda_{De}$, $0.705\lambda_{De}$, $1.785\lambda_{De}$, and $3.405\lambda_{De}$ in the top left, top right, bottom left, and bottom right plots, respectively. 71

3.8	Collisions involving two electrons each, illustrating that a small change of the electron starting positions leads to a large change in their positions at later times. The electrons in the collision shown by the solid lines start at $(-10, 0.1)$ and $(10, -0.1)$ while those in the collision shown by the dotted lines start at $(-10, 0.2)$ and $(10, -0.2)$.	72
3.9	The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in a BEPS simulation using the same parameters as the ddcMD simulation in Figure 3.7. The z position of the test charge is $0.705\lambda_{De}$ in the left plot and $3.405\lambda_{De}$ in the right plot.	73
3.10	The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in a BEPS simulation with approximately 152 cells per particle. The z position of the test charge is $3.94\lambda_{De}$ in the left plot and $7.42\lambda_{De}$ in the right plot. . .	74
3.11	The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in a BEPS simulation with 512 cells per particle. The z position of the test charge is $0.705\lambda_{De}$ in the left plot and $3.405\lambda_{De}$ in the right plot.	75
4.1	The integrands of the so-called wake (blue) and collisional (red) portions of Bohr stopping power as a function of impact parameter. Starting around $b \sim L_t$, the full dielectric integrand includes collective screening, and will fall below the collisional result, thereby removing the logarithmic large- b divergence.	93

4.2 Energy vs. distance traveled in ddcMD simulations using as projectiles anti-zinc nuclei (top), anti-neon nuclei (middle), and anti-protons (bottom). The dashed-dotted black line is the energy calculated using the Bohr formula (Equation 4.63), while the solid lines are the 27 simulation projectiles. 100

4.3 Energy vs. distance traveled in BEPS simulations using as projectiles anti-zinc nuclei in the top two plots and anti-protons in the bottom plot. The cell-width is λ_{De} in the top plot and $\lambda_{De}/8$ in the bottom two plots. The dashed-dotted black line is the stopping power calculated using the wake script while the solid lines are the 27 simulation projectiles. 101

4.4 The stopping powers measured in ddcMD (blue circles) and BEPS (red circles) simulations for anti-protons (left), anti-neon nuclei (middle), and anti-zinc nuclei (right). The ddcMD results are plotted with a cell width of zero for convenience. Projectiles in ddcMD start with $V_0 = 11.28v_{th}$ while those in BEPS start with $V_0 = 11.25v_{th}$. Stopping powers calculated using the Bohr formula and the wake script (red squares) are plotted for comparison with the ddcMD and BEPS results, respectively. Bohr 1 (blue squares) uses Equation 4.63 while Bohr 2 (blue diamonds) uses Equation 4.5, with the term inside the logarithm in Bohr 1 being twice that of Bohr 2. 103

4.5	The stopping powers measured in ddcMD and BEPS simulations for anti-zinc nuclei starting at various speeds along with comparisons to the Bohr stopping power from Equation 4.63 and the wake script, respectively. The cell width in BEPS is $0.125\lambda_{De}$. The ddcMD stopping powers are blue circles while the Bohr stopping power is a solid blue line. The BEPS stopping powers are red circles while the stopping power from the wake script is a solid red line.	104
4.6	Energy vs. distance traveled for electrons in ddcMD and BEPS simulations are in the top and bottom plots, respectively. The electrons start with $V_0 = 11.28v_{th}$ and $11.25v_{th}$ in ddcMD and BEPS, respectively. The BEPS simulation uses a cell width of $0.125\lambda_{De}$. The dashed-dotted black line is the energy calculated using the Bohr formula and the wake script for comparison with ddcMD and BEPS, respectively.	106
4.7	Perpendicular vs. parallel distance traveled for electrons in ddcMD and BEPS are in the top and bottom plots, respectively, for electrons starting with $V_0 = 11.28v_{th}$ and $11.25v_{th}$, respectively. The BEPS simulation uses a cell width of $0.125\lambda_{De}$. $s_{ }$ is the distance traveled along the electron's initial trajectory, while s_{\perp} is the distance traveled perpendicular to it.	107

4.8 Stopping power and angular displacement of electrons in ddcMD (blue circles) and BEPS (red circles) simulations. The ddcMD results are plotted with a cell width of zero for convenience. Projectiles in ddcMD start with $V_0 = 11.28v_{th}$ while those in BEPS start with $V_0 = 11.25v_{th}$. Stopping powers calculated using the Bohr formula and the wake script (red squares) are plotted for comparison with the ddcMD and BEPS results, respectively. Bohr 1 (blue squares) uses Equation 4.63 while Bohr 2 (blue diamonds) uses Equation 4.5, with the term inside the logarithm in Bohr 1 being twice that of Bohr 2. 108

4.9 The E_z field of an electron's wake at $t = 20\omega_{pe}^{-1}$ in an OSIRIS simulation with $\Delta = \Delta_0/2$ using the Yee solver. The slice is taken down the middle of the box in y . Numerical Čerenkov radiation, caused by the electron exceeding the propagation speed of light on the mesh, is radiated off to the sides. 111

4.10 The stopping powers measured in OSIRIS simulations using $\Delta = \Delta_0/4$ with the traditional Yee solver and the fourth-order accurate extended stencil. The red line is the stopping power in a simulation with a plasma, while the blue line is from a simulation without a plasma. 113

4.11 The stopping power in the QuickPIC simulations (blue with circles) compared with the fluid wake stopping power from Equation 4.30 (black with squares), the Bethe formula in Equation 4.6 (dashed green), the QED formula in Equation 4.75 (dashed red), and the Bohr formula in Equation 4.63 (dashed megenta). 115

5.1	The E_z field (left) and focusing field (right) of an electron wake in a QuickPIC simulation. The top row shows 2D slices of the fields, and the green dot represents the electron's position. Lineouts are taken along the black lines and plotted as red curves in the bottom row, whereas the fluid theory results are also plotted as dashed blue curves.	121
5.2	The focusing fields in the QuickPIC simulation. 2D slices of E_r and B_θ are on the left and right, respectively. The green dots represent the electron's position.	122
5.3	2D slices of E_z , the focusing field, and the density fluctuation on the left, center, and right, respectively. The wake acquires an upward curvature as we look farther behind the particle. The green dots represent the electron's position.	123
5.4	The E_z field of a wake in the OSIRIS simulation is on the left. The blue spot at the front of the wake is the electron's location. On the right is a comparison of the E_z field lineouts of the wake from QuickPIC in dashed-dotted green, from OSIRIS in dashed red, and from fluid theory in solid blue (for a point particle but off-axis by $\Delta_0/2$). The lineouts are taken along the electron's trajectory. . . .	124
5.5	Tracks of plasma particles in the wake in an OSIRIS simulation with $\Delta = \Delta_0$. The tracks are plotted in the frame of the moving window, and are in the same plane as the slices in Figure 5.3. The tracks have been shifted close to each other for ease of viewing. The x -axis is at the top of the box, to indicate the initial x -positions of the particles before they are affected by the test charge. The Displacement* axis at the bottom of the plot should be used as a scale for measuring a particle's displacement from its initial position at the top of the plot.	126

5.6	For the second track to the right of center in Figure 5.5, starting at the peak of the oscillation around $z - ct = 16\frac{c}{\omega_{pe}}$, the simulation track is in blue, while the numerical integration of Equation 5.6 is in red.	128
5.7	2D slices of E_z of a wake in OSIRIS for $10eV$, $100eV$, and $1keV$ plasma temperatures on the left, center, and right, respectively. The upward curvature turns downward as the plasma heats up. .	129
6.1	The Fast Ignition concept. After the compression stage, a high-intensity short-pulse laser hits the pellet of the compressed fuel, creating a beam of energetic electrons. As the electrons slow down in the core, they heat it, starting the ignition process.	131
6.2	Betatron motion of an electron in a wake as calculated using QuickPIC (yellow line) and a MATLAB script (black line). The electron starts at the blue dot and its tracks are plotted over the focusing field.	136
6.3	Betatron motion of five independent electrons in a QuickPIC wake as calculated using the MATLAB script. The tracks are imposed over the longitudinal field on the left and over the focusing field on the right. The bright green dots near the top of the plots represent the position of the electron producing the wake.	137
6.4	The path of a beam electron during a complete oscillation cycle, returning to its starting point, inside the wake produced by another beam electron, is on the left. The change in F_z during the cycle experienced by the electron due to its presence inside the wake is plotted on the right.	138

6.5	The E_z fields at $t = 100/\omega_{pe}$ for the 1keV OSIRIS simulation used to study beam electron motion in a thermal plasma. The subtracted field is on the left, showing the wakes, with the blue dots at the start of the wakes representing the beam electron positions. The unsubtracted field is on the right.	140
6.6	The motion of one beam electron in another's wake in OSIRIS simulations with various background temperatures. The transverse separation vs. time is on the left and the longitudinal separation on the right. The (blue, green, red, black) line corresponds to a simulation with a background temperature of (0eV, 10eV, 100eV, 1keV).	140
6.7	The number of beam electrons in a cubic skin depth for a range of beam and plasma densities. For reference, the red ticks on the x -axis are the plasma electron densities for a $450\text{g}/\text{cm}^3$ core with $1\mu\text{m}$ (1ω) and $0.5\mu\text{m}$ (2ω) lasers. The (linestyle, N_b , laser) triplets are (solid white, 1, 1ω), (dash-dotted red, 64, 1ω), (dashed green, 1, 2ω), and (dotted black, 64, 2ω).	141
6.8	The 125 lattice beam electrons at $t = 5/\omega_{pe}$ with $N_b = 1$	151
6.9	The relative forces on the 125 lattice beam electrons at $t = 5/\omega_{pe}$ with $N_b = 1$ are on the left. On the right, the forces are plotted over the E_z field in a slice across the middle of the box in y	151
6.10	The relative forces on the 125 beam electrons for the lattice case with $N_b = 1$ are plotted on the left and right at $t = 3,005/\omega_{pe}$ and $6,895/\omega_{pe}$, respectively.	152

6.11	The stopping power for many different electron configurations, all with $N_b = 1$, averaged over all the beam electrons and plotted as a multiple of the single-electron stopping power. The horizontal dashed black line is the single-electron stopping power. The dashed purple line is the lattice case, and all other lines are for simulations where the electrons are initialized with random deviations from their lattice centers. The shaded areas represent the standard deviation for lines of the same color. On the right, there are 16 solid lines, corresponding to 16 different simulations.	152
6.12	The electron beam configuration at $t = 5/\omega_{pe}$ with the electrons displaced random distances from their lattice centers.	153
6.13	A comparison of the correlated stopping power between $\Delta = \Delta_0$ and $\Delta_0/2$. with a decreased cell width. The stopping power is plotted as a multiple of the single-electron stopping power when $\Delta = \Delta_0/2$. The (dashed blue, dashed-dotted green) and (solid red, dotted magenta) lines are the lattice and random placement cases for a cell width of $(\Delta_0, \Delta_0/2)$, respectively. The shaded areas represent the standard deviation for lines of the same color. The horizontal (dash-dotted gray, dashed black) line is the single-electron stopping power for $\Delta = (\Delta_0, \Delta_0/2)$	154
6.14	For various initial electron beam densities using 125 electrons, the average stopping power for the lattice cases is plotted on the left, and the stopping power averaged across all 16 random placement runs is plotted on the right. The (line style, N_b) pairs are (dash-dotted yellow, 1/64), (dashed black, 1/27), (solid blue, 1/8), (dotted green, 1), (dash-dotted red, 8), (dashed cyan, 64), (solid magenta, 1,000). The shaded areas represent the standard deviation for lines of the same color.	156

6.15	For various initial electron beam densities using 125 electrons, the average peak stopping power is plotted as a multiple of the single-electron stopping power measured in QuickPIC for the lattice cases in dashed red and random placement cases in solid blue. The modified multiple of the Bohr stopping power is plotted for the lattice cases in dash-dotted magenta and for the random placement cases in dotted black.	158
6.16	For various electron beam densities initially inside $1,000 (c/\omega_{pe})^3$, the average stopping power for the lattice cases is plotted on the left, and the stopping power averaged across all 16 random placement runs is plotted on the right. The (line style, N_b) pairs are (solid blue, 1/8), (dotted green, 1), (dash-dotted red, 8), and (dashed cyan, 64). The shaded areas represent the standard deviation for lines of the same color.	160
6.17	For various electron beam densities initially inside $1,000 (c/\omega_{pe})^3$, the average peak stopping power is plotted as a multiple of the single-electron stopping power measured in QuickPIC for the lattice cases in dashed red and random placement cases in solid blue. The modified multiple of the Bohr stopping power is plotted for the lattice cases in dash-dotted magenta and for the random placement cases in dotted black.	161
6.18	The electron positions are plotted on the left in 3D and the beam density is plotted on the right for a 2D slice through the middle of the box in y . The plots are at $t = 600/\omega_{pe}$, and the run is one of the random electron placement cases with $N_b = 64$ initially inside the $1,000 (c/\omega_{pe})^3$ cube.	164

6.19	The E_z field and the \hat{z} component of the plasma current for a slice through the middle of the box in y are plotted on the left and right, respectively, at $t = 600/\omega_{pe}$ for the simulation in Figure 6.18. . . .	164
6.20	The phase space density as a function of P_z and z at $t = 600/\omega_{pe}$ for the simulation in Figure 6.18.	165
6.21	For various electron beam densities in the box-length beams, the average stopping power for the lattice cases is plotted on the left, and the stopping power averaged across all eight random placement runs is plotted on the right. The (line style, N_b) pairs are (solid blue, 1/8), (dotted green, 1), (dash-dotted red, 8), and (dashed cyan, 64). The shaded areas represent the standard deviation for lines of the same color.	168
6.22	For various electron beam densities inside box-length beams, the average peak stopping power is plotted as a multiple of the single-electron stopping power measured in QuickPIC for the lattice cases in dashed red and random placement cases in solid blue. The modified multiple of the Bohr stopping power is plotted for the lattice cases in dash-dotted magenta and for the random placement cases in dotted black.	169
6.23	The transverse average of E_z and the lineout down the middle of the box along z at the times when peak average stopping power occurs. The beam electrons are randomly placed and $N_b = 64$. . .	170
6.24	The electron positions are plotted on the left in 3D and the beam density is plotted on the right for a 2D slice through the middle of the box in y . The plots are at $t = 140/\omega_{pe}$ for the same run as in Figure 6.23.	172

6.25	The E_z field and the plasma z current in a slice through the middle of the box through y are plotted on the left and right, respectively, at $t = 140/\omega_{pe}$ for the same simulation as in Figure 6.23.	173
6.26	The phase space density as a function of P_z and z at $t = 140/\omega_{pe}$ for the same simulation as in Figure 6.24.	174
A.1	A schematic of a 2D plasma sheet moving across the box in QuickPIC. As it crosses the box, it responds to the beam, producing the wake. The green dots are the beam electrons.	182
A.2	The 3D flow chart of QuickPIC with the 2D flow chart embedded inside it. The 3D flow chart is repeated for each s -step. During each s -step, the 2D routine is called, and the 2D cycle repeats for each ξ -step as the 2D plasma sheet crosses the box.	182

LIST OF TABLES

2.1	The bounce times of deeply trapped electrons measured in the simulations of Figure 2.14, along with the bounce times calculated using the average plasma wave field amplitude along the particle’s trajectory.	37
4.1	The parameters for the electrostatic single-particle stopping power simulations using ddcMD and BEPS.	97
4.2	The parameters for the OSIRIS stopping power simulations. . . .	110
4.3	The parameters for the QuickPIC stopping power simulations. . .	114
6.1	Baseline parameters for the QuickPIC correlated stopping simulations.	135
6.2	The peak stopping powers of a typical electron in box-length beams averaged over the eight random placement simulations for each value of N_b , along with the time-steps used in those simulations. .	166
6.3	The peak timing for the average stopping power from the random placement runs at each N_b . The secondary peaks to not appear during the simulations for $N_b = 1$ and $1/8$	171

ACKNOWLEDGMENTS

Thanks to David Strozzi for acting as an unofficial adviser during my time Lawrence Livermore National Laboratory (LLNL), for showing me what a good place it is to do research as a student during the summer of 2009, and for proposing the ideas for my thesis research.

Thanks go to Warren Mori for overseeing my research and allowing me to head off to LLNL to perform my research rather than staying at UCLA.

Thank you to Frank Graziani for being my technical supervisor at LLNL. Without his support, I would not have been able to perform my research at LLNL and would not have attended so many conferences. He had lots of patience with my research, even as it morphed into something different to the original plan.

Thanks to all the administrators who helped arrange everything for me, and to Christine Zachow in particular, who oversees the scholar and fellowship programs at LLNL, for working with me to arrange postdoc events and help me navigate issues at LLNL.

Thanks to Jim Glosli, David Richards, Mike Surh, Bruce Langdon, Michael Murillo, Paul Grabowski, Denise Hinkel, Ed Williams, Frank Tsung, Viktor Decyk, Ben Winjum, Weiming An, Thomas Grismayer, and Jay Fahlen for their assistance with my thesis research. Apologies to those individuals who I forgot to list.

Thanks to David Michta for being a good friend during my time at LLNL and for assistance with difficult mathematical and conceptual problems.

Thanks go to Jorge Rocca and his then graduate student Mark Berrill for introducing me to the field of plasma physics during my undergraduate career. I worked in Jorge's lab for a couple years performing plasma physics research. They taught me about basic plasma processes and helped me cultivate programming skills in Matlab, which proved very useful throughout my graduate career.

Finally, thanks go to Adam Tableman for occasional programming and computer assistance and providing amusement whenever I returned to UCLA with his unique if sometimes cynical perspective on science and life in general.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and by the University of California, Los Angeles under Grants FG52-09NA29552 and DE-NA0001833. This work was funded in part by the Laboratory Directed Research and Development Program at LLNL under project tracking codes 08-ERD-017, 09-SI-011, 11-SI-002, and 12-SI-005. Work by Ian N. Ellis was supported in part by the Lawrence Scholar Program at LLNL.

VITA

- 2004-2006 Undergraduate Student Researcher for Professor Jorge Rocca
NSF Engineering Research Center for Extreme Ultraviolet Science and Technology
Department of Electrical Engineering
Colorado State University
- 2006 B.S. Physics
Colorado State University
- 2007-2009 Teaching Assistant
Department of Physics & Astronomy
University of California, Los Angeles
- Summer 2008 Graduate Student Researcher
Plasma Simulation Group
Department of Physics & Astronomy
University of California, Los Angeles
- 2009 M.S. Physics
University of California, Los Angeles
- Summer 2009 Temporary Student Intern
Lawrence Livermore National Laboratory
- 2010-2014 Lawrence Scholar
Lawrence Livermore National Laboratory
- 2014-2017 Future Technical Leader
Northrop Grumman Information Systems

PUBLICATIONS

1. Ian N. Ellis, Frank R. Graziani, James N. Glosli, David J. Strozzi, Michael P. Surh, David F. Richards, Viktor K. Decyk, and Warren B. Mori. “Studies of particle wake potentials in plasmas.” *High Energy Density Physics*, **7**(3):191-196, Sep 2011.
2. Michael P. Surh, Ian N. Ellis, James N. Glosli, Frank R. Graziani, William D. Krauss, Michael S. Murillo, David F. Richards, and Fred H. Streitz. “Large-Scale Molecular Dynamics Simulation of Charged Particle Energy Deposition in Plasmas.” *IEEE Transactions on Plasma Science*, **39**(11):2620-2621, Nov 2011.
3. Ian N. Ellis, David J. Strozzi, Benjamin J. Winjum, Frank S. Tsung, Thomas Grismayer, Warren B. Mori, Jay E. Fahlen, and Edward A. Williams. “Convective Raman amplification of light pulses causing kinetic inflation in inertial fusion plasmas.” *Physics of Plasmas*, **19**(11):112704, 2012.
4. Josh May, John Tonge, Ian N. Ellis, Warren B. Mori, Frederico Fiúza, Ricardo A. Fonseca, Louís O. Silva, and Chuang Ren. “Enhanced Stopping of Macro-Particles in Particle- in-Cell Simulations.” *Physics of Plasmas*. Submitted.
5. Ian N. Ellis, David J. Strozzi, Warren B. Mori, Viktor K. Decyk, Frank R. Graziani, and Frank S. Tsung. “Correlated electron stopping in Fast Ignition plasmas.” *Physical Review Letters*. Working paper.
6. Michael S. Murillo, Ian N. Ellis, Paul E. Grabowski, Michael P. Surh, David F. Richards, James N. Glosli, A. Bruce Langdon, and Frank R. Graziani.

“Stopping power in dense plasmas using molecular dynamics.” *Physical Review E*. Working paper.

CHAPTER 1

Introduction

In this dissertation, we study the importance of plasma wave wakes. Electron plasma waves have been studied extensively since the work of Langmuir, Bohm, and Gross. Much of the linear and nonlinear theory of plasma waves, including kinetic treatments, make the traveling wave assumption and watch the evolution of the wave in time. However, plasma waves are often excited by moving drivers and evolve in space and time. Such plasma waves can be viewed as wakes left behind by electromagnetic waves or charged particles in high energy density plasmas. Electromagnetic waves commonly produce wakes in plasmas via three-wave decay processes [Kru03, DKL74] such as stimulated Raman and [FKL75b, FKL75a], the parametric decay instability [DG67, KD69, SW72], and two-plasmon decay [LR76, YRL12]. Charged particles produce wakes as they move through a plasma [Dec87, KT73]. The studies in this dissertation focus on how plasma wave wakes affect stimulated Raman scattering and electron beam stopping [BD08, DF99, SB08, ASD09], which play important roles in Inertial Confinement Fusion (ICF) [AV09, LAB04].

The wakes that appear in our stimulated Raman scattering and electron beam stopping research take the form of electron plasma waves (EPWs), which are commonly referred to as Langmuir waves [Lan28, TL29]. No ion waves appear in the studies at hand as the simulations use frozen ions or a smooth neutralizing background to simplify the study and suppress processes such as the Langmuir decay instability, in which an EPW decays into an ion acoustic wave and another

EPW [DG67, Gol84, DFL00].

The work on Raman scattering focuses exclusively on backward stimulated Raman scattering (BSRS), in which an incident light wave in a plasma decays into a backward-propagating light wave and a forward-propagating plasma wave. BSRS can be a significant source of drive energy loss in ICF because it scatters light away from the target. It can also degrade the uniformity of the drive in indirect drive if the energy loss is different between the inner and outer beams. Furthermore, the plasma waves can generate superthermal electrons, which can preheat the core and/or the ablator. The study focuses on BSRS in ICF-relevant laser and plasma conditions using OSIRIS [FST02] particle-in-cell (PIC) [Daw83, BL85] simulations. The study has three distinct parts. The first part is a detailed comparison of the linear amplification of a well-defined seed pulse propagating counter to the pump laser using coupled-mode theory and OSIRIS simulations. To compare theory to OSIRIS simulations, finite-size particle corrections are included in the theory. This part of the study demonstrates excellent agreement between the coupled mode description and the OSIRIS simulations and demonstrates that special relativity corrections affect the seed amplification for ICF parameters. The second part explores how this seed pulse can trigger large reflectivities after it leaves the simulated plasma. The seed creates an EPW in the plasma, which the pump laser continues to drive and which traps electrons, reducing the Landau damping and leading to high laser reflectivities through a process known as kinetic inflation [VDB01, VDB02, VDB07, SWL07]. The final part of the study examines the onset of inflationary scattering and the bursty nature of BSRS using continuous seeds. Non-resonant seeds, which have frequencies not at the peak of the linear gain curve, require higher incident intensity to cause inflation. When the seed frequency is non-resonant, the reflected light is modulated with a period inversely proportional to the difference between the seed and resonant frequencies. This three-part study provides many details about linear amplification via BSRS

and the effects of seeds in BSRS overall.

The study of electron beam stopping is motivated by Fast Ignition (FI) [THG94, AV09, ASB07, STL12]. In this ICF concept, the target is compressed to a dense state, then a high-intensity short-pulse laser is fired at a point on the surface of the target, generating a beam of electrons which travels to the core and heats it to ignition temperatures. The heating of the core depends on the electron stopping power. The stopping power of a single electron, in turn, depends in large part on the plasma wave wake made by the electron. The wake contributes to the stopping power since the energy left behind in the wake comes from the particle, or equivalently, the electric field of the wake evaluated at the particle slows it down. In addition, electrons propagating behind one electron can interact with its wake and may also enhance it. We study this possible coherent enhancement of wakes. We call electron beam stopping from the enhancement of wakes “correlated stopping.” There is agreement on the single-electron stopping power [ASD09, SB08, LP06, SD05], but correlation effects have been largely overlooked, aside from Refs. [BD08, DF99, MF02]. This study demonstrates that correlation effects usually lead to an increase in stopping power, and the stopping power of the electron beam is likely much greater than previously thought.

Due to the importance of electron wakes in correlated stopping, the study begins with a comparison of the wakes in PIC simulations with current models of wakes. In the non-relativistic regime, the wake in electrostatic PIC simulations agrees extremely well with a calculation using Vlasov theory. In the relativistic regime, cold fluid theory agrees reasonably with the wavelength of the wake observed in OSIRIS and QuickPIC [HDR06, ADM13] simulations, along with the form of the wake during the first wavelength, but misses a transverse spreading phenomenon which we attribute to the fact that the plasma electron oscillations are preferentially in the radial direction. As shown by Dawson, such oscillations have a frequency that depends on the ratio of the displacement to its initial radius.

The agreement between the simulations and theories, and the physical explanation for the differences, provide confidence that our simulation tools can be used to study correlated stopping.

Various stopping power models are briefly discussed, including the Bohr and Bethe models of classical and quantum mechanical stopping, and a model that includes quantum electrodynamics. Comparisons are made with PIC and molecular dynamics (MD) simulations. Stopping power for fast ($v \gg v_{th}$) heavy ions in ddcMD [RGC09, GBB12], which is an MD code that uses the particle-particle particle-mesh (PPPM) method [HGE73], agrees very well with the Bohr formula. Simulations of the stopping power in electrostatic PIC produces stopping power below that from Bohr formula, but which agrees with calculations using Vlasov theory taking into account finite-size particles and field smoothing. As the particle (cell size) is reduced, the PIC results gradually converge to the ddcMD results. Simulations of fast electrons in both MD and electrostatic PIC exhibit a significant spread in energy with time.

Stopping power of relativistic electrons does not simply depend on Coulomb interactions. Therefore, neither electrostatic PIC nor ddcMD can be used. Instead, we resort to electromagnetic PIC. However, simulations of electrons with momentum of $10m_e c$ in OSIRIS exhibit significant numerical Čerenkov radiation at small cell widths, which overwhelms the stopping power due to the plasma, though recent work suggests that a new solver eliminates this problem. Numerical Čerenkov arises because the phase velocity of light on the grid is less than the speed of light, allowing particles traveling near the speed of light to exceed the phase velocity of light. The same simulations are performed with QuickPIC, whose model is free of numerical Čerenkov radiation. They demonstrate good agreement with a fluid model of stopping power using the cell-width as a cut-off until the stopping power in the simulations saturates when a cell-width is a quarter of the inter-particle spacing. These results demonstrate the benefits and

limitations of various types of simulations when used to study stopping power.

Correlated stopping simulations of mono-energetic, mono-directional electron beams with electron momentum of $10m_e c$ in a cold FI core plasma are performed. These simulations use one simulation electron for each real electron and an electron size equal to the plasma electron interparticle spacing but much less than the beam interparticle spacing. All the simulations with beam interparticle spacings of $2c/\omega_{pe}$ or less show an increase in stopping power over that of a single electron. Time-dynamics of the beams plays a fundamental role because the beam-plasma-like instability [DMO70, GR71] and self-focusing [RSC89, RSC90] increase the stopping power with time. We say “like” because, for the parameters used here, the interparticle spacing can be larger than the wavelength and width of the wake.

Three distinct varieties of beam stopping simulations are performed. First, simulations with 125 beam electrons initialized in a cube are carried out, which demonstrate that the stopping power gradually increases as the spacing, and hence the beam size, decrease due to bunching of the beam electrons. Second, the stopping power of various beam densities initialized in a cube with a volume of $1,000 (c/\omega_{pe})^3$ is examined, demonstrating that correlated stopping depends strongly on beam density. In these simulations, the stopping power increases first via transverse beam-plasma-like instabilities, such as filamentation and self-focusing, then via the longitudinal beam-plasma-like instability. Finally, beams that are longer in the propagation direction are studied using various beam densities. These final beams have dimensions of $10c/\omega_{pe} \times 10c/\omega_{pe} \times 80c/\omega_{pe}$. These simulations show spatio-temporal growth (growth with propagation distance and growth from the head to the tail of the beam) that is consistent with streaming instabilities. When the beam electron separation is $0.25c/\omega_{pe}$, the stopping of a typical beam electron peaks at $(1 \pm 3) \times 10^3$ times that of an uncorrelated electron. These three varieties of beam stopping simulations demonstrate that correlation effects increase the stopping power—sometimes dramatically—under certain conditions and exhibit

a complex time-dependence. Such increases imply that more beam energy might be transferred into the core of an FI target than previously thought. However, the beams simulated represent a very small volume of what is expected in Fast Ignition and the beams have no energy spread or transverse divergence (emittance). Therefore, this work also points to directions for future work such as studying larger beams, beams with large transverse divergence (emittance) and energy spreads, as well as plasmas with higher temperatures.

This dissertation addresses in detail two examples of the complex role that wakes play in ICF plasmas. BSRS illustrates one of the many processes through which wakes can be problematic and hinder progress towards ignition. Correlated stopping points towards a relatively unexplored path along which wakes can help achieve ignition. The results from both areas of research have important implications for ICF.

CHAPTER 2

Convective Raman Amplification of Light Pulses Causing Kinetic Inflation in Inertial Fusion Plasmas

2.1 Introduction

Backward stimulated Raman scattering [DKL74, FKL75b, FKL75a] (BSRS) in plasmas is a process in which an incident light wave in a plasma decays into a backward-propagating light wave and a forward-propagating plasma wave. During this process, the plasma wave is excited in a finite region of space and time. It is a driven wave packet (wake) during the onset of the instability, and a freely propagating wake after saturation. In this sense, SRS is relevant to the topic of this dissertation.

BSRS has been a subject of much study, in large part because it scatters light away from the target in inertial confinement fusion (ICF) [AV09, LAB04]. Early research focused on relatively high intensities, where growth was in the weakly damped convective and absolutely unstable regimes, and saturation occurred due to wave-breaking and/or pump depletion [KA75, EKH89]. The competition between back-, forward-, and side-scatter was also investigated [EK83, KEL80]. As laser and plasma parameters for ICF evolved, research in BSRS shifted to the strongly damped regime. In modern experiments, SRS typically occurs at densities and temperatures for which $k\lambda_{De} \gtrsim 0.3$, where Landau damping is signifi-

cant. In this kinetic regime, the measured BSRS reflectivities can greatly exceed the values from linear theory calculations, e.g. in the single-hot-spot experiments of Reference [MCF02]. Although 3D paraxial-envelope simulations with linear damping have correctly modeled the intensity threshold for SRS in experiments with a smoothed, multi-speckle beam in a pre-formed uniform hohlraum plasma [FDL09], such models do not include much nonlinear and kinetic physics. For example, a process called kinetic inflation was proposed to explain the single-hot-spot results [VDB01, VDB02, VDB07, SWL07]. In kinetic inflation, a small-amplitude plasma wave excited in the strongly damped convectively unstable regime can trap electrons, modifying the distribution function so that the kinetic damping of the plasma wave is greatly reduced or vanishes [ON65, MO72]. Therefore, for the same incident laser intensity, SRS can then transition to the weakly damped or absolutely unstable regime [Win10]. There has also been recent work on how BSRS in the kinetic regime can saturate due to nonlinear frequency shifts [VDB01, VDB02, WFT10] or related trapped-particle instabilities [BV04] caused by electron trapping. Recent research has demonstrated the importance of the propagation and evolution of plasma wave packets, including how the reflected light can occur in bursts spaced proportionally to the inverse of the nonlinear frequency shift [WFT10]. Recently, it has been shown that hot electrons and back-and side-scattered SRS produced by one speckle interact with neighboring speckles [Win12]. Very recently, simulations of many speckles have shown that SRS can self-organize between speckles, producing coherent bursts of SRS [YAR12a]. These simulations are difficult to analyze and little work has explored in detail how scattered light, plasma waves, or the resulting changes to the electron distribution in one region of space or time enhancing SRS at different times or locations.

In this Chapter, we make a detailed comparison of the linear amplification of a well defined counter-propagating seed pulse using coupled-mode theory and OSIRIS PIC [FST02] simulations. We then explore how this seed pulse can trigger

large reflectivities after it has left the plasma using OSIRIS simulations. We consider situations in which no (little) BSRS occurs with only the pump (no seed is used). The seed intensity and pulse length are varied. For short seed pulses ($\sim 500\omega_0^{-1}$ FWHM), the seed pulse is linearly amplified as it transits the box. This amplification agrees with linear theory when appropriately modified to take into account special relativity and the use of finite-size particles in PIC codes [Daw83, BL85], such as OSIRIS. We simulate ICF-relevant laser and plasma conditions, and demonstrate that, even for plasma temperatures less than 5 keV, special relativity increases the linear gain and shifts down the scattered light wavelength. We find, for these short seed pulses, that kinetic inflation occurs after the seed pulse leaves the box. The timing and amplitude of the first peak in reflectivity depends on the duration and intensity of the seed pulse. We examine the trajectories of trapped particles to verify that kinetic inflation is occurring and that the bounce period is consistent with the amplitude of the Langmuir wave (which we also call the plasma wave).

For longer seed pulses, the inflationary burst of scattered light overlaps with the seed. Under these conditions, the measured gain of the seed can reach several times the steady-state linear gain value when the seed wavelength is near the peak of the gain spectrum. We also examine the onset of inflationary scattering and the bursty nature of BSRS using continuous seeds. We also vary the seed frequency and find that non-resonant seeds, which are not at the peak of the linear gain curve, require higher incident intensity to cause inflation. We also observe that, when the seed frequency is non-resonant, the reflected light is modulated with a period inversely proportional to the difference between the seed and resonant frequencies. (In this Chapter, we use the term “resonance,” where $1 + \chi_r = 0$ for electrostatic waves, and “peak gain” interchangeably.) After inflation sets in, the measured gain of the seed *decreases* with incident seed intensity due to pump depletion.




Figure 2.1: The geometry of the OSIRIS simulations.

The Chapter is outlined as follows. We present the simulation geometry and plasma conditions in Section 2.2, and discuss the linear theory of convective BSRS gain and its relativistic and PIC modifications in Section 2.3. We describe in Section 2.4 a subtraction technique that we use in our data analysis. In Section 2.5, we discuss the amplification of short-duration seed pulses in PIC simulations and as calculated by a coupled-mode solver. Section 2.6 covers our observation of kinetic inflation that occurs after the seed pulse passes, and Section 2.7 covers our measurements of kinetic inflation using longer duration seeds. Finally, we discuss the onset of inflation with continuous seeds in Section 2.8 and conclude in Section 2.9.

2.2 Geometry and Plasma Conditions

Throughout this Chapter, we use normalized units to describe the OSIRIS simulations. To make connection to parameters of interest for ICF, we assume that the incident pump has a wavelength of $\lambda_0 = 351\text{nm}$. When we note quantities in physical units, they correspond to this pump wavelength. Our formulas and other quantities are given in CGS units while the temperature is often given in units of eV.

Figure 2.1 depicts the usual simulation geometry. The pump laser (ω_0, \vec{k}_0)

is incident from the left with a normalized electric field amplitude $eE_0/m_e c \omega_0 = 3.68 \times 10^{-3}$ (in physical units, $I_0 = 1.5 \times 10^{14}$ W/cm²). The seed (ω_1, \vec{k}_1) enters the simulation box from the right and beats with the pump, inducing an electrostatic wave (ω_2, \vec{k}_2) in the plasma, which travels to the right. Light from the pump scatters off the plasma wave and amplifies the seed. For the simulation parameters, this process is a convective instability [Str05].

We use a fixed ion neutralizing background and we do not add collisions. We simulate densities and temperatures for which SRS is expected to occur for ignition experiments at the National Ignition Facility (NIF) [MBR09]. The plasma has a uniform density $n = 0.12n_c$ ($\approx 1.1 \times 10^{21}$ cm⁻³ for $\lambda_0 = 351$ nm), where n_c is the critical density for the pump laser. The electron thermal speed is $v_{th} = 0.0699c$ ($T_e = 2.5$ keV). The box has length $L = 1790c/\omega_0$ (100 μ m) and 8192 cells, giving a cell width of $1.1\lambda_D$, where λ_D is the Debye length. We use 16,384 particles per cell and quadratic splines for the particle shape to reduce spurious grid heating and noise due to aliasing. The particle boundary conditions are thermalizing for particles and we use perfectly matched layers [JNG05, BHP08, Ber94] for the field boundary conditions. The particle pusher is relativistically correct [BL85] in all the simulations except where we state otherwise.

For ICF plasma conditions in thermodynamic equilibrium, the primary instigator of BSRS is electromagnetic noise caused by Thomson scattering [SWH08]. The seed intensities we use in our simulations are far above the level of this background noise. For the plasma conditions used in our simulations and a typical NIF quad (four laser beams arranged in a square) with effective F-number [MBR09] of 8 as the pump, the effective Thomson scattering seed within the FWHM of the peak relativistic gain (discussed below) is 1.6×10^4 W/cm² = $1.1 \times 10^{-10}I_0$. SRS growth from such noise is initially linear, and enhanced over plane-wave growth by intense speckles in a phase-plate-smoothed beam. We note that it is not possible to directly measure the noise on NIF, and it may be higher than that from

Thomson scattering. Convective amplification of the noise can create scattered light amplitudes far above the thermal noise, and may transition to regimes where kinetic effects are significant. We choose seed intensities to induce such effects.

2.3 Linear Theory of Convective BSRS Gain

2.3.1 Non-Relativistic Theory

Our PIC simulations never reach a steady state. However, linear theory states that the reflectivity will quickly reach a steady state when we use a continuous scattered light wave seed. In the convective steady state, the seed intensity is amplified by a factor of e^G by the time it exits the box. G is the linear intensity gain exponent, commonly called “the gain,” and we present an equation for it in the strong damping limit. Here, we summarize the results from a detailed derivation for the steady-state gain from Reference [SWH08].

Given the pump (ω_0, \vec{k}_0) and the seed (ω_1, \vec{k}_1) , we calculate the plasma wave (ω_2, \vec{k}_2) using the matching conditions,

$$\omega_0 = \omega_1 + \omega_2 \tag{2.1a}$$

and

$$\vec{k}_0 = \vec{k}_1 + \vec{k}_2. \tag{2.1b}$$

We need ω_2 and \vec{k}_2 for evaluating the plasma susceptibility when we calculate the gain.

Let the seed intensity be denoted by $I_1(z)$, where $z = (0, L)$ is the (left, right) edge of the box. Then,

$$I_1(z) = e^{G_l(z)} I_1(L), \tag{2.2}$$

where $G_l(z)$ is the linear intensity gain exponent,

$$G_l(z) \equiv \int_z^L \Gamma_1(z') I_0(z') dz'. \tag{2.3}$$

We neglect pump depletion and light wave damping, so I_0 is constant, which leads to

$$\Gamma_1 \equiv \Gamma_s \text{Im} \left[\frac{\chi_e}{\varepsilon} (1 + \chi_I) \right], \quad (2.4)$$

where the subscripts e and I denote the electron and ion species, respectively, χ_j is the collisionless susceptibility for species j , $\varepsilon(k_2, \omega_2) = 1 + \sum_j \chi_j(k_2, \omega_2)$ is the plasma dielectric function, and

$$\Gamma_s \equiv \frac{2\pi r_e}{m_e c^2} \frac{1}{\omega_0} \frac{k_2^2}{k_0 |k_1|}, \quad (2.5)$$

where $r_e \equiv e^2/m_e c^2$ is the classical electron radius. For a Maxwellian velocity distribution, χ_j is given by

$$\chi_j(k_2, \omega_2) = -\frac{\omega_{pj}^2}{2k_2^2 v_{Tj}^2} Z' \left(\frac{\omega_2}{\sqrt{2} k_2 v_{Tj}} \right), \quad (2.6)$$

where ω_{pj} is the plasma frequency of species j , $v_{Tj} = \sqrt{T_j/m_j}$ is the thermal speed of species j , and $Z'(s) \equiv dZ/ds$. $Z'(s)$ must be calculated numerically, and is typically found by first computing $Z(s)$, the plasma dispersion function [FC61]. It can be shown that $Z'(s) = -2sZ(s) - 2$, with

$$Z(s) = i\sqrt{\pi} e^{-s^2} [1 + \text{erf}(is)]. \quad (2.7)$$

$\chi_j \rightarrow -(\omega_{pj}/\omega_2)^2$ as $m_j \rightarrow \infty$, so we can set $\chi_I \rightarrow 0$ everywhere (recall we use fixed ions in the OSIRIS simulations). In particular, $\varepsilon = 1 + \chi_e$ and

$$\Gamma_1 = \Gamma_s \text{Im} \left[\frac{\chi_e}{\varepsilon} \right]. \quad (2.8)$$

We further simplify Equation 2.3 since we are dealing with a uniform plasma. Γ_1 is constant, so the gain is given by

$$G_l(z) = \Gamma_1 I_0 (L - z). \quad (2.9)$$

We also define an amplitude gain rate, g_0 , as

$$g_0 = \frac{\Gamma_1 I_0}{2} \propto \frac{\chi_i}{(1 + \chi_r)^2 + \chi_i^2}, \quad (2.10)$$

where χ_r and χ_i are the real and imaginary parts of χ , respectively. We plot the theoretical gain spectrum for the conditions of interest in Figure 2.2 as dash-dotted lines.

Equations 2.2-2.5 are valid in the strong damping limit. This limit applies when $|v_{g2}\partial a_2/\partial x| \ll |\nu_2 a_2|$, where v_{g2} is the plasma wave group velocity, a_2 is the plasma wave action amplitude (defined in Section 2.5.2), and ν_2 is the plasma wave damping rate. For regimes of interest, ν_2 is dominated by Landau damping [Lan46]. In a homogeneous plasma in the convective steady-state, which is where our gain calculation applies, this condition is $g_0 \ll \nu_2/v_{g2}$. Working at the peak of the non-relativistic gain curve, we have a spatial gain rate of $g_0^{nr} = 3.28 \times 10^{-4} \omega_0/c$ and $\nu_2/v_{g2} = 0.0611 \omega_0/c$. Therefore, at NIF, we are in the strong damping limit.

2.3.2 Relativistic Modification

We now explore the impact of special relativity on linear gain. Estabrook and Kruer [EK83] included an analysis of SRS for temperatures for which relativistic corrections are important, and performed 1.5D PIC simulations of laser and plasma conditions where the plasma wave is weakly damped. They found that non-relativistic linear theory does not adequately describe the wavenumber of the fastest growing mode in high-temperature (~ 64 keV) plasmas, but taking into account the effective (reduced) plasma frequency and corresponding density due to special relativity brings theory and simulation into better agreement. More recently, Bergman and Eliasson derived a fully relativistic expression for the unmagnetized plasma dielectric function [BE01], and Bers et. al. derived approximate expressions relevant to current and near-future deuterium-tritium fusion plasmas [BSS09]. Palastro et. al. have also derived a fully relativistic description of Thomson scattering [PRP10].

In this subsection, we simply make some heuristic changes to the formulas in

the previous subsection to account for special relativity. We replace the susceptibility in Equation 2.8 with the relativistic one of Bergman and Eliasson, which is computed using a 3D Jüttner-Synge distribution, as opposed to a Maxwellian distribution.

Given the 3D Jüttner-Synge distribution,

$$f(\gamma) = \frac{n_0 \mu e^{-\mu\gamma}}{4\pi m_e^3 c^3 K_2(\mu)}, \quad (2.11)$$

the electron susceptibility is given by

$$\chi_e(\kappa_2, \Omega_2) = \frac{\mu}{\kappa_2^2} \left[1 - \frac{\mu}{K_2(\mu)} \frac{\partial^2}{\partial \mu^2} \frac{P(\mu, \kappa_2/\Omega_2)}{\mu} \right], \quad (2.12)$$

where γ is the relativistic Lorentz factor, $\mu \equiv m_e c^2 / T_e$, $K_2(\mu)$ is the modified Bessel function of the second kind, $\Omega_2 \equiv \omega_2 / \omega_{pe}$, and $\kappa_2 \equiv k_2 c / \omega_{pe}$. P is given by

$$P(\mu, \kappa_2/\Omega_2) \equiv \int_1^\infty \frac{e^{-\mu\gamma}}{\sqrt{\gamma^2 - 1}} \frac{d\gamma}{\gamma^2 (1 - \kappa_2^2/\Omega_2^2) + \kappa_2^2/\Omega_2^2} - \frac{i\pi\sigma}{2\kappa_2/\Omega_2} e^{-\mu(\kappa_2/\Omega_2)/\sqrt{\kappa_2^2/\Omega_2^2 - 1}}, \quad (2.13)$$

where $\sigma = 0$ for $\Omega_2^2 \geq \kappa_2^2$ and $\sigma = 1$ for $\Omega_2^2 < \kappa_2^2$.

We can also take into account relativistic effects in the electromagnetic dispersion relation, $\omega^2 = \omega_{pe}^2 + c^2 k^2$, by using a relativistic version of the plasma frequency:

$$\begin{aligned} \omega_{pe}^2 &\rightarrow \omega_{pe}^2 \frac{\mu^2}{K_2(\mu)} \int_1^\infty \frac{\partial^2}{\partial \mu^2} \left(\frac{e^{-\mu\gamma}}{\mu} \right) \frac{\sqrt{\gamma^2 - 1}}{\gamma^4} d\gamma \\ &\approx \omega_{pe}^2 \left(1 - \frac{5}{2\mu} \right) \quad \text{for } \mu \gg 1. \end{aligned} \quad (2.14)$$

We use this relativistic plasma frequency when we calculate k_0 and k_1 . The decrease in the effective plasma frequency with temperature is due to relativistic corrections to the internal energy of the plasma [TM98]. However, this change only has a relatively small impact on the gain at the temperature we use in our simulations. It is also possible to calculate k_0 and k_1 using the fully relativistic

transverse dispersion relation from Reference [BE01], but we did not attempt to do so.

The overall effect of special relativity on the gain curve is shown in Figure 2.2. Notice that the peak of the relativistic gain curve lies above the peak of the non-relativistic gain curve, and the peak occurs at a shorter wavelength. This difference occurs because there are less particles and a shallower slope at the plasma wave’s phase velocity in the Jüttner distribution than in the Maxwellian distribution, resulting in weaker Landau damping. At the peak of the analytic non-relativistic gain curve ($\lambda_1 = 1.659\lambda_0 = 582.37$ nm), for which $k_2\lambda_{De} = 0.289$, the non-relativistic damping rate is $3.24 \times 10^{-3}\omega_0$, while the relativistic rate is $2.44 \times 10^{-3}\omega_0$. The strong damping limit still applies at the peak of the relativistic gain curve ($\lambda_1 = 1.655\lambda_0 = 580.88$ nm), for which the spatial gain rate is $g_0^r = 4.43 \times 10^{-4}\omega_0/c$ and $\nu_2/v_{g2} = 0.0465\omega_0/c$.

We also performed some gain calculations using the approximate expression of Bers et. al. for the relativistic longitudinal dispersion relation. Their expression shifts the gain curve down in wavelength significantly more than numerically integrating the formula of Bergman and Eliasson.

2.3.3 PIC Modification

We can improve the agreement between simulation and theoretical results by taking into account a few known aspects of finite-difference PIC codes: finite-size particles, differencing operators, and field smoothing plus compensation [BL85]. The particles have a finite size because the charge and current are interpolated to a grid via the “shape factor” $S(\vec{x})$:

$$q\delta(\vec{x}) \rightarrow qS(\vec{x}). \tag{2.15}$$

To reduce the self-heating and spurious noise from aliasing, we often use second-order B-splines. Transforming to Fourier space,

$$qS(\vec{x}) \rightarrow qS(\vec{k}), \quad (2.16)$$

where

$$S(k) = \left(\frac{\sin(k\Delta/2)}{k\Delta/2} \right)^3 \quad (2.17)$$

for 1D simulations, with Δ being the cell width.

In finite-difference codes, like OSIRIS, differencing operators modify the dispersion relation by changing the relationship between the charge density ρ , longitudinal electric field E_2 , and electrostatic potential ϕ . The finite-difference representation for the derivative is

$$\frac{\phi_{j+1/2} - \phi_{j-1/2}}{\Delta}, \quad (2.18)$$

which, in Fourier space, is

$$k_2 \frac{\sin(k_2\Delta/2)}{k_2\Delta/2} \equiv K(k_2) \quad (2.19)$$

. Therefore, in Fourier space, Poisson's equation becomes

$$\begin{aligned} 4\pi\rho(k_2) &= k_2^2 \left(\frac{\sin(k_2\Delta/2)}{k_2\Delta/2} \right)^2 \phi(k_2) \\ &= K^2(k_2)\phi(k_2) \end{aligned} \quad (2.20)$$

and

$$\begin{aligned} E_2(k_2) &= -ik_2 \frac{\sin(k_2\Delta)}{k_2\Delta} \phi(k_2) \\ &= -iK(k_2)\phi(k_2). \end{aligned} \quad (2.21)$$

We additionally smooth the fields in our simulations to further reduce the effects of aliasing, and we compensate to reduce numerical modifications to the dispersion relation for small \vec{k} . Without the use of splines and smoothing, grid heating instabilities occur for $\Delta \gtrsim 3\lambda_{De}$ [BL85]. In OSIRIS, we use both splines

and smoothing to eliminate grid heating. We perform the filtering (smoothing or compensation) of some quantity ϕ on the grid by replacing

$$\phi_j \text{ with } \frac{W\phi_{j-1} + \phi_j + W\phi_{j+1}}{1 + 2W}, \quad (2.22)$$

where j is the grid index and W is a weighting factor. Transforming into Fourier space, this operator becomes

$$\begin{aligned} \phi_f(k) &= \frac{1 + 2W \cos(k\Delta)}{1 + 2W} \phi_0(k) \\ &= SM_W(k\Delta) \phi_0(k). \end{aligned} \quad (2.23)$$

We use two types of filters in the simulations presented in this Chapter. The first is a two-pass filter, which we use unless stated otherwise. The first pass has a stencil of $\frac{1}{4}(1,2,1)$ ($W=1/2$) and the second has a stencil of $\frac{1}{4}(-1,6,-1)$ ($W=-1/6$). Therefore,

$$\begin{aligned} SM(k) &= SM_{1/2}(k) SM_{-1/6}(k) \\ &= \frac{1 + \cos(k\Delta)}{2} \frac{3 - \cos(k\Delta)}{2}. \end{aligned} \quad (2.24)$$

The second filter has five passes, and we choose it because it causes less deviation from the longitudinal dispersion relation without quadratic interpolation effects than the 2-pass filter. The first four passes use a stencil of $\frac{1}{4}(1,2,1)$ ($W=1/2$) and the last pass uses a stencil of $\frac{1}{4}(-5,14,-5)$ ($W=-5/14$). We note that compensation for higher-order particle shapes can be obtained. We perform additional simulations with this filter to observe the effect on the gain curve.

$$\begin{aligned} SM(k) &= SM_{1/2}^4(k) SM_{-5/14}(k) \\ &= \left(\frac{1 + \cos(k\Delta)}{2} \right)^4 \frac{14 - 10 \cos(k\Delta)}{4}. \end{aligned} \quad (2.25)$$

The particle shape factor, differencing operators, and field smoothing affect the plasma frequency through the relationship

$$\omega_{pe}^2 \rightarrow \omega_{pe}^2 \frac{k_2 \kappa(k_2)}{K^2(k_2)} (L \cdot S(k_2))^2 SM(k_2). \quad (2.26)$$




Figure 2.2: Several calculations of the linear convective BRS gain spectra for a box $1790c/\omega_0$ ($100 \mu\text{m}$) long. The dash-dotted curves are analytic results, and the (dashed, solid) ones take into account PIC effects with the (5,2)-pass filter.

We make this transformation everywhere ω_{pe} appears in our gain formulas to account for PIC (finite size particle and filtering) effects.

The effect of the shape factor, differencing operators, and both the 2-pass and 5-pass filters on the relativistic and non-relativistic gain curves is shown in Figure 2.2. Checking the strong damping limit condition for the 2-pass filter, in the non-relativistic case, the peak gain rate drops to $g_0^{nr} = 2.75 \times 10^{-4} \omega_0/c$ with $\nu_2/v_{g2} = 0.0636 \omega_0/c$, while for the relativistic case, it drops to $g_0^r = 3.70 \times 10^{-4} \omega_0/c$ with $\nu_2/v_{g2} = 0.0531 \omega_0/c$.

2.4 Subtraction Technique

We use a subtraction technique [Dec87] in our data analysis to clearly see waves with amplitudes below the background plasma fluctuation level. The technique requires running two simulations, the first with a perturbation whose effects we wish to examine, and the second without the perturbation, but with the same random number generator seed. We then subtract the results of the second simulation from the results of the first. In our case, the first simulation has both a backward propagating light seed pulse and a forward propagating pump, while the second simulation has just the pump. We use the subtraction technique to produce all the plots of transverse and longitudinal fields unless stated otherwise.

Figure 2.3 shows the electrostatic field induced by the beating of the pump and the seed pulse. The seed pulse in this simulation has a Gaussian-like profile, as described in section 2.5, with $\lambda_1 = 1.644\lambda_0$ (577nm) and $I_{1s} = 5 \times 10^{-4}I_0$. The amplitude of the plasma wave is so small that we cannot distinguish it from the background fluctuations without using the subtraction technique. In the subtracted result, background fluctuations enter the simulation starting at the sides of the box due to the thermalizing boundary conditions. Fluctuations will always re-enter the subtracted data as the two simulations become uncorrelated. However, the thermalizing boundaries re-emit particles using a random number generator, and a single generator is used for all the particles on a given processor. Once the order of the particles reaching the boundary changes, fluctuations rapidly enter the subtracted data.

Because we do not observe meaningful SRS without a seed, we also use the subtraction technique as a means of separating the scattered light from the pump light when pump depletion is not significant. When pump depletion is small, the subtraction technique for the transverse electric field works well for finding the scattered light at all positions in the box. However, when pump depletion




Figure 2.3: The longitudinal electric field in a simulation with (a) and without (b) a seed pulse. Subtracting the latter from the former reveals the plasma wave (c,d).

becomes significant, the subtraction technique does not produce good results by itself anywhere except at the far left side of the box, where pump depletion does not occur. Therefore, we can still use the subtraction technique to observe the reflectivity at the left side of the box, but we need to filter out the pump in Fourier space to observe the scattered light anywhere else in the box.

The number of particles per cell in the simulation affects the usefulness of the subtraction technique. Fluctuations take longer to enter the simulation as we increase the number of particles per cell. When we decrease the number of particles per cell, low-intensity seed pulses become more difficult to distinguish from the background fluctuations when we use the subtraction technique, until we cannot distinguish the peak of a pulse with $I_{1s} = 5 \times 10^{-4} I_0$ from the fluctuations in simulations with 512 particles per cell. However, changing the number of particles per cell has no significant effect on the convective amplification of the seed pulse.

2.5 Convective Gain Simulations & Coupled-Mode Results

2.5.1 OSIRIS Convective Gain Simulations

We perform OSIRIS PIC simulations to observe seed amplification in the linear regime and determine under what conditions SRS enters the nonlinear regime. The simulations begin at $t = 0$ with the pump incident from the left. After the pump has fully propagated across the plasma, $2,000\omega_0^{-1}$, the seed enters from the right. We use two different temporal profiles, or “shapes,” for the seed amplitude and vary its wavelength and intensity. The first shape is approximately Gaussian, and rises from zero to its peak amplitude in $\tau = 500\omega_0^{-1}$, then falls back to zero over another $500\omega_0^{-1}$, for $\sim 500\omega_0^{-1}$ FWHM. The second shape is a flat-top pulse with a Gaussian-like rise and fall time of $\tau = 200\omega_0^{-1}$, and a constant peak




Figure 2.4: The shapes of the seed pulses in the simulations. A Gaussian-like pulse (a) with a $500\omega_0^{-1}$ rise and fall time ($\sim 500\omega_0^{-1}$ FWHM), and a flat-top pulse (b) with a $200\omega_0^{-1}$ Gaussian-like rise and fall, and a steady amplitude for $600\omega_0^{-1}$ in between.

amplitude for $600\omega_0^{-1}$ in between. The two pulse shapes are plotted in Figure 2.4. We describe the seed pulse using the notation $I_1(z = L, t) = I_{1s}s(t)$, so that $s(t)$ describes the pulse shape and I_{1s} is the maximum incident intensity. In our simulations using seed pulses, $I_{1s} \geq 5 \times 10^{-4}I_0$.

Figure 2.5 shows the scattered light and the plasma wave as a function of position and time in an OSIRIS PIC simulation when we use a Gaussian-like seed pulse with $I_{1s} = 5 \times 10^{-4}I_0$ and $\lambda_1 = 1.644\lambda_0$. We also include line-outs of the scattered light amplitude vs. position at various times to show the evolution of the seed pulse more clearly as it crosses the box from right to left. We use a Hilbert transform to envelope the results, producing a smooth appearance.

We define the measured gain, g_{meas} , of a pulse as

$$g_{meas} \equiv \ln \left(\frac{\max(I_1(z = 0, t))}{I_{1s}} \right). \quad (2.27)$$

The plot at the top of Figure 2.6 shows the gain we measure in simulations when we vary the seed wavelength while keeping I_{1s} fixed at $5 \times 10^{-4}I_0$. The linear rel-




Figure 2.5: Scattered light vs. position and time (a), lineouts of the scattered light vs. position at various times (b), and plasma wave amplitude (c) from a simulation using a Gaussian-like seed pulse with $I_{1s} = 5 \times 10^{-4} I_0$ and $\lambda_1 = 1.644\lambda_0$.




Figure 2.6: Measured gain in simulations with box length $1790c/\omega_0$ ($100 \mu\text{m}$) vs. seed wavelength (a) and seed intensity I_{1s} (b). Plotted are the gains we measure using a Gaussian-like pulse (green with square markers) and a flat-top pulse (blue with circle markers), both with $\lambda_1 = 1.644\lambda_0$. (a) includes the gain from simulations with a Gaussian-like pulse and the 5-pass filter (black with 'x' markers), plus those without relativistic effects (brown with '+' markers). Several theoretical gain curves taking into account PIC effects are included: 2-pass non-relativistic is dash-dotted red, 2-pass relativistic is solid magenta, and 5-pass relativistic is dashed purple.

ativistic gain peaks near $\lambda_1 = 1.644\lambda_0$ (577nm) in simulations with a 2-pass filter and near $\lambda_1 = 1.650\lambda_0$ (579nm) in simulations with a 5-pass filter. Our simulation results agree with these predictions. We also plot a gain curve from simulations using a non-relativistic particle pusher and a Maxwellian velocity distribution. This non-relativistic curve lies below the relativistic one at most points, as we expect, and the location of its peak agrees with the gain curve from non-relativistic theory with a 2-pass filter.

In the plot on the bottom of Figure 2.6, we see how the measured gain changes as we vary I_{1s} while keeping the seed wavelength fixed at $\lambda_1 = 1.644\lambda_0$. The measured gain of the Gaussian-like pulse remains relatively constant as we increase the initial seed intensity, until the seed intensity reaches several times the pump intensity. This behavior indicates that we are in the linear regime for even large seed amplitudes. The deviation from linear theory at the highest amplitudes is due to pump depletion. However, the measured gain of the flat-top pulse increases with seed intensity before falling off. This slight increase in the gain is caused by the onset of inflation where the higher amplitude seed pulse generates a larger amplitude plasma wave. Trapped particles execute several bounces, thus decreasing the Landau damping rate while the seed is still present. We explain this effect in more detail in the following sections. We note that if we had decreased (increased) the seed pulse length, the deviation from linear behavior would occur at higher (lower) seed intensity.

2.5.2 The Coupled-Mode Equations

The measured gain of a seed pulse can differ from the steady-state linear result due to several linear and nonlinear effects. Linear effects include pulse shape, with each frequency in a spectrum of incident frequencies being amplified at a different rate. Nonlinear effects include pump depletion and nonlinear (kinetic) and nonlocal reductions to the real part of the frequency and the damping rate of

the plasma wave. In this section, we investigate the effect of the pulse shape on the measured gain using the coupled-mode equations [BSW98, Str05]. A comparison of the coupled-mode and OSIRIS results for NIF conditions, isolates linear from truly nonlinear, kinetic physics, and provides confidence in the PIC method. We note that Wang et. al. compared OSIRIS against reduced models for lower plasma temperatures, which are of interest in Raman amplification experiments [WML09].

In the coupled-mode equations, we let (ω_i, \vec{k}_i) of the carrier waves be real and work with complex envelopes $a_i(\vec{x}, t)$. We assume that the envelopes vary slowly with respect to the carriers, such that $|\nabla a_i| \ll |\vec{k}_i a_i|$ and $|\partial a_i / \partial t| \ll |\omega_i a_i|$. The complex envelopes for the action amplitudes a_j relate to the physical quantities by

$$\vec{A}_j = -i \left(\frac{2\pi}{\omega_j} \right)^{1/2} a_j \exp[i(\vec{k}_j \cdot \vec{x} - \omega_j t)] \hat{k} + cc, \quad j = 0, 1 \quad (2.28)$$

for light waves, with A_j the vector potential, and by

$$n_1 = \frac{ik_2}{2} \left(\frac{2n_B}{m_e \omega_2} \right)^{1/2} a_2 \exp[i(\vec{k}_2 \cdot \vec{x} - \omega_2 t)] + cc, \quad (2.29)$$

for the plasma wave. n_B is the spatially varying background electron density, n_1 is the perturbation on top of it, and cc denotes complex conjugate.

The couple mode equations are

$$\left(\frac{\partial}{\partial t} + \vec{v}_{g0} \cdot \nabla + \nu_0 + i\delta_0 \right) a_0 = K a_1 a_2, \quad (2.30)$$

$$\left(\frac{\partial}{\partial t} + \vec{v}_{g1} \cdot \nabla + \nu_1 + i\delta_1 \right) a_1 = -K a_0 a_2^*, \quad (2.31)$$

$$\left(\frac{\partial}{\partial t} + \vec{v}_{g2} \cdot \nabla + \nu_2 + i\delta_2 \right) a_2 = -K a_0 a_1^*, \quad (2.32)$$

where the coupling constant is

$$K \equiv \frac{k_2}{\sqrt{\omega_0 \omega_1 \omega_2}} \frac{\omega_{pe}^2}{\sqrt{8n_B m_e}}, \quad (2.33)$$

ν_i is the damping rate of mode i , and, for light waves, $\delta_{0,1} \equiv (\omega_{pe}^2 - \omega_{0,1}^2 + c^2 k_{0,1}^2) / (2\omega_{0,1})$ is the detuning of the mode from its natural frequency. We neglect




Figure 2.7: The BRS reflected light using a flat-top seed with $\lambda_1 = 1.644\lambda_0$ (a) and $1.650\lambda_0$ (b). The simulation results are plotted with solid magenta lines while the coupled-mode results are plotted with dash-dotted black lines. The horizontal dashed red line indicates the maximum seed amplitude with no gain.

inverse bremsstrahlung, which is negligible in the PIC simulations, so $\nu_0 = \nu_1 = 0$. We assume nu_2 is due to Landau damping, $\nu_2 = \varepsilon_i / (\partial\varepsilon_r / \partial\omega_2)$, where $\varepsilon_r = \text{Re}[\varepsilon]$, $\varepsilon_i = \text{Im}[\varepsilon]$, and ε is the kinetic dielectric function. We set the light-wave detuning, $\delta_0 = \delta_1 = 0$. For the electrostatic mode, we find δ_2 using the kinetic equation, $\delta_2 = -\varepsilon_r / (\partial\varepsilon_r / \partial\omega_2)$. The equations in Sections 2.3.2 and 2.3.3 allow us to take into account special relativity and PIC effects when we calculate the coefficients.

We can directly compare the results from OSIRIS and the coupled-mode solver by examining the reflected light. Figure 2.7 shows the reflected light from runs with $I_{1s} = 5 \times 10^{-4} I_0$ using flat-top pulses with $\lambda_1 = 1.644\lambda_0$ and $1.650\lambda_0$. In Figure 2.8, we compare the measured gain from PIC simulations and the coupled-mode solver for various wavelengths. For these quasi-linear cases, the simulation and coupled-mode results are in excellent agreement for both the Gaussian-like and flat-top pulse runs.




Figure 2.8: The measured gain in a box of length $1790c/\omega_0$ ($100 \mu\text{m}$) predicted by the coupled-mode solver as we vary the wavelength using a Gaussian-like pulse (dashed brown with ‘+’ markers) and a flat-top pulse (dashed magenta with circle markers). For comparison, we include the PIC simulation result for the Gaussian-like pulse (solid blue with ‘x’ markers) and the flat-top pulse (solid red with square markers) along with the linear relativistic gain curve taking into account PIC effects (dash-dotted green).




Figure 2.9: The longitudinal field vs. space at $t = 4,000\omega_0^{-1}$ when using a flat-top seed with $\lambda_1 = 1.644\lambda_0$ (a) and $1.650\lambda_0$ (b); the same runs as used in Figure 2.7. The simulation results are plotted with solid magenta lines while the coupled-mode results are plotted with dash-dotted black lines.

The agreement between simulation and coupled-mode results is not as good when we examine the longitudinal field. Figure 2.9 shows the amplitude of the longitudinal field at $t = 4,000\omega_0^{-1}$ for the same runs as shown in Figure 2.7. The disagreement occurs soon after the wave begins growing and is visible at about $x = 300c/\omega_0$. We are not yet sure of the reason for this disagreement. The disagreement becomes worse as background fluctuations begin to enter the PIC results after $x = 400c/\omega_0$. The disagreement for large x may be due more to the limitation of the subtraction technique than the physics.

2.6 Inflation After Seed Passage

For short seed pulses, we find significant reflectivity well after the seed propagates out of the simulation box for pump intensities which do not produce BSRS on their own. The pump now directly interacts with the Langmuir wave that is still

present in the plasma after the seed leaves the box. Without trapped particles, the Langmuir wave is described by its linear dispersion relation, BSRS remains in the strongly-damped limit, and no observable growth of BSRS occurs for the pump intensity and the plasma length of interest. However, a small amplitude wave can evolve into a nonlinear, weakly damped wave after the trapped particles execute a few bounces [ON65, MO72]. The period for a bounce, or bounce time, is $\tau_B \equiv 2\pi/\omega_B$, where $\omega_B = \sqrt{eE_2k/m_e}$ is the bounce frequency for deeply trapped electrons, E_2 is the electric field amplitude, and k is the wavenumber of the wave. As particles are trapped, the damping rate decreases below its linear Landau damping value to a residual level which depends on details of the problem [BMG12]. The ponderomotive beating of the pump and the scattered light will drive the wave to increasing amplitudes. Such a situation will lead to noticeable reflectivity later in the simulation. The seed pulse must be the cause of any such reflectivity because, with the pump amplitude we use in our simulations, BSRS is negligible without a seed. BSRS that occurs after the seed has left the box is both useful for isolating the process of kinetic inflation [VDB01, VDB02] and is potentially relevant to situations where BSRS in one region of space or time seeds BSRS in another one, creating a plasma wave that triggers an inflationary burst of BSRS.

Figure 2.10 shows the scattered light and plasma wave for times after the seed leaves the box in a simulation using a Gaussian-like seed pulse with $I_{1s} = 8 \times 10^{-3} I_0$ and $\lambda_1 = 1.644\lambda_0$. We observe high reflectivity after $t = 10,000\omega_0^{-1}$, along with a corresponding growth in the plasma wave. Figure 2.11 shows lineouts of the scattered light and longitudinal field at $x = 550c/\omega_0$ along with a comparison with the coupled-mode result. For the longitudinal field from the PIC simulation, we filter out all modes except those in the range $1.4\omega_0/c \leq k \leq 1.5\omega_0/c$. Notice the dip in the plasma wave amplitude around $t = 4,500\omega_0^{-1}$ in Figure Figure 2.11, corresponding to the drop in the seed's amplitude, before the plasma wave begins




Figure 2.10: The scattered light (a) and the plasma wave (b) seen when we extend the duration of the simulation using a Gaussian-like seed pulse with $I_{1s} = 8 \times 10^{-3} I_0$ and $\lambda_1 = 1.644 \lambda_0$. The seed exits the box around $t = 5,000 \omega_0^{-1}$.

to grow again.

We verify that kinetic inflation is occurring by tracking particles traveling near the plasma wave phase velocity and plotting their orbits in the frame of the wave, as done in Figure 2.12 from $t = 5,000\omega_0^{-1}$ to $t = 7,000\omega_0^{-1}$. The elliptic trajectories are clear indicators of particle trapping. During this time, the distribution function begins to nonlinearly flatten around the Langmuir wave phase velocity, as seen in Figure 2.13, which is another indication of particle trapping and an indication a nonlinear mode with a reduced damping rate. The tail is flattened over a wider range of velocities, which extends predominately to higher velocities, during the larger burst of SRS that grows after the seed leaves; for example, as shown at $t = 17,000\omega_0^{-1}$. This larger tail in our simulations is either the population of energetic electrons (commonly referred to as “hot electrons”) seen in recent experiments [DTH10] or a source of electrons that can be further accelerated in other plasma waves [WFT13]. The production of hot electron tails by SRS is an active area of research [WFT13, YAR12b], but the small flattening early in time is sufficient to affect the growth and further onset of SRS we are studying here.

Since particle trapping causes the inflationary scattering at late times, the inflationary bursts will occur earlier if the seed drives a larger Langmuir wave with a shorter τ_B , so that the trapped particles accumulate bounces faster. We can increase the Langmuir wave amplitude by increasing the intensity, the duration, or choosing a seed wavelength that produces a higher gain.

We first vary the initial amplitude of the seed. Figure 2.14 shows that the first burst of reflected light has a maximum around $t = 17,000\omega_0^{-1}$ when we use a Gaussian-like seed pulse with $I_{1s} = 4 \times 10^{-3}I_0$ and $\lambda_1 = 1.644\lambda_0$. As we increase I_{1s} , the burst moves earlier, but the difference in amplitude between the burst and final seed amplitude also decreases. As a clear demonstration of the effect of the higher intensity seeds, in Table 2.1 we examine the plasma wave amplitudes and particle bounce times for the simulations in Figure 2.14. We




Figure 2.11: Evolution of the scattered light (a) and the longitudinal field (the PIC simulation field is filtered) (b) with time at $x = 550c/\omega_0$ for the simulation of Figure 2.10. The simulation results are plotted with solid blue lines and the coupled-mode results are plotted with dashed black lines.




Figure 2.12: Orbits of particles trapped in the plasma wave in the simulation of Figure 2.10. The orbits are plotted for $2,000\omega_0^{-1}$ starting at $t_0 = 5,000\omega_0^{-1}$ in the wave frame, where $v_{ph} = 0.271c$ is the plasma wave phase speed and the orbits are centered around the relativistic $p_{ph} = 0.282m_e c$.




Figure 2.13: The distribution functions from the simulation of Figure 2.10 at $t = 6,000\omega_0^{-1}$ (solid blue) and $t = 17,000\omega_0^{-1}$ (dashed black) along with the Jüttner distribution for a 2.5keV electron plasma (dash-dotted red). The measured distributions have a plateau beginning around $\gamma\beta = 0.26$, indicating particle trapping.

I_{1s}	Avg. EPW Amp.	Calc. τ_B	Meas. τ_B
$[I_0]$	$[E_0]$	$[\omega_0^{-1}]$	$[\omega_0^{-1}]$
4×10^{-3}	0.020	610	675
8×10^{-3}	0.033	480	550
0.016	0.049	390	415
0.032	0.071	320	350

Table 2.1: The bounce times of deeply trapped electrons measured in the simulations of Figure 2.14, along with the bounce times calculated using the average plasma wave field amplitude along the particle’s trajectory.

examine a selection of particles near the plasma wave phase velocity between $x = 350c/\omega_0$ and $x = 450c/\omega_0$ at $t = 5,000\omega_0^{-1}$. We use the most deeply trapped particle to determine the time it takes to complete one bounce. To calculate the theoretical bounce time, we filter out all plasma wave modes except those in the range $1.4\omega_0/c \leq k \leq 1.5\omega_0/c$ and average the field amplitude over the bounce time along the particle’s track. We then substitute the measured values of the amplitude and wavenumber ($k = 1.44\omega_0/c$) into the formula for the bounce time. The results are in Table 2.1. The measured bounce times are slightly ($\sim 10\%$) longer than the simple expression. We believe this discrepancy is due to the fact that the amplitude of the wave and its phase velocity is changing with time, and because the calculation uses a linear approximation to the sinusoidal potential well.

Besides lowering the kinetic damping rate, another well-known effect of particle trapping is the nonlinear frequency down-shift of the plasma wave [MO72]. As the plasma wave grows, it will shift downward in frequency because it will trap more particles. According to the frequency matching condition in Equation 2.1a, the down-shift in the frequency of the plasma wave should be accompanied by an up-shift in the frequency of the scattered light. We examine this down-shift using




Figure 2.14: The reflected light in extended duration simulations using a Gaussian-like seed pulse with $\lambda_1 = 1.644\lambda_0$ and various maximum initial intensities. Shown are $I_{1s} = 4 \times 10^{-3}I_0$ (solid blue), $8 \times 10^{-3}I_0$ (dotted green), $0.016I_0$ (dash-dotted red), and $0.032I_0$ (dashed black).

a Wigner transform with a Choi filter [Mal99]. The Wigner transform takes a function of time and computes its representation as a function of both frequency and time. It maps $f(t) \rightarrow f(\omega, t)$.

The Wigner transform results for a run with a flat-top seed with $I_{1s} = 8 \times 10^{-3}I_0$ and $\lambda_1 = 1.638\lambda_0$ are shown in Figure 2.15. The results for an identical run, except with a different seed wavelength, $\lambda_1 = 1.627\lambda_0$, are shown in Figure 2.16. In both cases, the seed appears in the Wigner transform scattered light plots around $t = 4500\omega_0^{-1}$, when the seed reaches the left side of the box. The scattered light frequency shifts up while the plasma wave frequency shifts down, as expected due to trapping.

The inflationary bursts begin growing near the frequency with the highest gain regardless of the seed's central frequency, consistent with a harmonic oscillator that is driven off-resonance. The initial growth is near the central frequency of the seed in the case with $\lambda_1 = 1.638\lambda_0$, and in the seed's lower-frequency tail in the case with $\lambda_1 = 1.627\lambda_0$. This observation suggests that the ponderomotive beating of the seed and the pump simply disturbs the plasma and provides an initial level for growth, but the subsequent growth occurs at the most unstable mode.

There are two inflationary bursts in the scattered light plot in Figure 2.15. According to Reference [WFT10], the separation of the bursts in time is $\sim 2\pi/\Delta\omega_{NL}$, where $\Delta\omega_{NL}$ is the nonlinear shift of the scattered light and plasma wave from their non-inflationary resonant frequencies, where peak linear gain occurs. The bursts result from the pump and nonlinearly shifted reflected light driving a plasma wave off resonance in a region where the plasma is unperturbed. A harmonic oscillator driver off resonance drives a plasma wave at the average frequency, which grows and decays with a period given by $\sim 2\pi/\Delta\omega$. Similar behavior is seen here. However, we see the bursts occur at different frequencies. The first is roughly on-resonance while the second has $\Delta\omega_{NL} \approx 0.002\omega_0$. Averaging the two, we have




Figure 2.15: The scattered light (a) and plasma wave (b) in a simulation using a flat-top seed pulse with $I_{1s} = 8 \times 10^{-3} I_0$ and $\lambda_1 = 1.638 \lambda_0$. Below them are the Wigner transforms of the reflected light at $x = 0$ (c), the plasma wave at $x = 500c/\omega_0$ (d), and the plasma wave at $x = 900c/\omega_0$ (e).




Figure 2.16: The scattered light (a) and plasma wave (b) in a simulation using a flat-top seed pulse with $I_{1s} = 8 \times 10^{-3} I_0$ and $\lambda_1 = 1.627 \lambda_0$. Below them are the Wigner transforms of the reflected light at $x = 0$ (c) and the plasma wave at $x = 900c/\omega_0$ (d).

$\Delta\omega_{NL}^{avg} = 0.001\omega_0$, or a separation of about $6,000\omega_0^{-1}$, which is in good agreement with the actual separation of the two bursts. Figures 2.15 and 2.16 show the seed wavelength also affects the time required for inflation to set in. We explore this effect in Sections 2.7 and 2.8.

Increasing the width of the seed drives the plasma wave longer, increasing its amplitude and causing inflationary bursts to occur earlier. In Figure 2.17, we vary the duration of a flat-top pulse while keeping the rise and fall time at $200\omega_0^{-1}$. The burst of high reflectivity moves earlier and, correspondingly, the plasma wave reaches a high amplitude quicker as we increase the pulse duration from $1,000\omega_0^{-1}$ to $3,000\omega_0^{-1}$. However, this effect only occurs if we drive near resonance. If we drive off resonance, then the large bursts don't occur, and the reflectivity oscillates with a period of $2\pi/\Delta\omega_{NR}$, where $\Delta\omega_{NR}$ is the difference between the seed frequency and the resonant frequency, at which inflationary SRS grows. This is completely explained from the response of a harmonic oscillator driven off resonance, as described above.

2.7 Inflation of the Seed

When the seed pulse width becomes comparable to the bounce time, then the seed itself can undergo inflation, illustrating the difference between the linear and inflationary regimes. We perform simulations using a flat-top pulse with a duration of $6,000\omega_0^{-1}$ with $I_{1s} = 8 \times 10^{-3}I_0$ at various wavelengths. The reflectivity plot in Figure 2.18 demonstrates the variation of the scattered light with time in the simulations as compared to the steady-state values from linear theory. The reflectivities seen in the simulations can significantly exceed the linear theory values, particularly for the seeds with $\lambda_1 = 1.638\lambda_0$ and $1.644\lambda_0$. The runs using seeds with $\lambda_1 = 1.632\lambda_0$ and $1.650\lambda_0$ reach levels above the linear theory values, but dip below linear values several times due to driving off resonance. In these four




Figure 2.17: Results from simulations with flat-top seed pulses with $I_{1s} = 8 \times 10^{-3} I_0$ and $\lambda_1 = 1.644 \lambda_0$. The scattered light (left) and plasma wave (right) for simulations using a pulse duration of $(1000, 3000)\omega_0^{-1}$ are in panels (a-b, c-d), respectively.




Figure 2.18: The reflected light seen in several simulations using a flat-top seed of duration $6,000\omega_0^{-1}$ with $I_{1s} = 8 \times 10^{-3}I_0$ for various wavelengths. The two top curves are for seeds with $\lambda_1 = 1.638\lambda_0$ (solid red) and $1.644\lambda_0$ (dashed cyan). The four lower curves are for seeds with $\lambda_1 = 1.627\lambda_0$ (dotted blue), $1.632\lambda_0$ (dashed green), $1.650\lambda_0$ (dash-dotted purple), and $1.658\lambda_0$ (solid yellow). For comparison, we mark the steady-state linear relativistic PIC values using horizontal dashes on the left side of the plot. The red vertical dash on the lower right side of the plot indicates approximately when the seeds end.

cases, inflationary scattering continues after the seed ends around $\omega_0 t = 10, 200$. When we use seeds with $\lambda_1 = 1.627\lambda_0$ and $1.658\lambda_0$, the reflectivity does not reach significantly above the linear value and drops down when the seed ends, because inflationary scattering does not occur easily when the seed is far from resonance.

We observe oscillations in the reflectivity in all the cases, except when we use the seeds with $\lambda_1 = 1.638\lambda_0$ and $1.644\lambda_0$. These oscillations are due to the scattered light seed driving SRS off resonance [WFT10]. The ponderomotive beating of the pump and scattered light drives the plasma wave, so if the phase of the beat drive leads the plasma wave by $\pi/2$, the plasma wave no longer grows. The plasma wave density $n_1 \propto -\frac{\partial}{\partial x} E_2$, while the beat drive $F_p \propto -\frac{\partial}{\partial x} E_0 E_1$. The product,

$$R_p = \left(\frac{\partial}{\partial x} E_2 \right) \left(\frac{\partial}{\partial x} E_0 E_1 \right), \quad (2.34)$$

indicates the phase difference between the beat drive and the plasma wave [Win10].

We also note that the change in the pump energy density with time is given by

$$\frac{\partial W_0}{\partial t} = -\frac{e}{8\pi m_e \omega_1 k_2} R_p. \quad (2.35)$$

Therefore, if R_p is positive, the waves are in phase, the beat drive is resonantly driving the plasma wave, and energy is transferred from the pump to the seed and plasma wave. The inverse applies if R_p is negative.

Figure 2.19 shows the scattered light and the resonance plot for the seed with $\lambda_1 = 1.644\lambda_0$. We smooth the result from the resonance diagnostic in x using a 6-point moving average. Notice the dip (valley) in the scattered light amplitude around $x = 1250c/\omega_0$, $t = 7,000\omega_0^{-1}$, and the corresponding negative area on the resonance plot around $x = 1500c/\omega_0$. When this drop occurs, the beat drive and the plasma wave are out of phase, so energy flows from the scattered light wave to the pump. This shift away from resonance is due to the nonlinear frequency shift of the plasma wave, explained in Section 2.6. The non-resonant drive is also responsible for the oscillations we see in the reflectivity plot of Figure 2.18.




Figure 2.19: The scattered light (a), and the resonance product R_p from Equation 2.34 (b) for the flat-top seed of duration $6,000\omega_0^{-1}$ with $\lambda_1 = 1.644\lambda_0$.

Figure 2.20 shows the same plots as Figure 2.19, except using a seed with $\lambda_1 = 1.638\lambda_0$. This seed continuously drives the scattered light to its maximum amplitude before dropping off, whereas the one with $\lambda_1 = 1.644\lambda_0$ drives a lower-amplitude burst of scattered light before it produces a second burst at much higher amplitude. Notice that the resonance plot in Figure 2.20 shows that the beat wave and plasma wave are in resonance until the peak of the scattered light burst. This resonant drive leads to a burst of scattered light at higher amplitude than in the $\lambda_1 = 1.644\lambda_0$ case.

2.8 Inflation of Continuous Seeds

In this section, we extend the duration of the flat-top seed pulse so that it remains on through the end of the simulation. Based on what we have observed in Sections 2.6 and 2.7, we expect the inflationary behavior we see in these simulations to depend on the seed intensity and wavelength. We know that there must be an intensity threshold below which the seed does not drive inflationary behavior in the simulations (very long bounce times), because we see negligible scattering without a seed. However, as we saw in Section 2.7, a seed that is intense enough to drive inflation on resonance may not be intense enough to drive it off resonance.

Furthermore, the use of continuous seeds is related to past work by Winjum *et. al.* on scattering off of plasma wave packets [WFT10]. Light scattering off a wave packet that has undergone a nonlinear frequency shift acts as a seed for BSRS in the unperturbed background plasma. However, the beat drive frequency is not a natural mode of the background plasma. This non-resonant beat drive leads to oscillations in the reflected light with a period $2\pi/\Delta\omega_{NL}$, where $\Delta\omega_{NL}$ is the nonlinear frequency shift of the plasma wave, much like the amplitude of a simple harmonic oscillator varies when driven off resonance. In this section, we examine the effect of resonant and non-resonant drive in more detail by using continuous




Figure 2.20: The scattered light (a), and the resonance product R_p from Equation 2.34 (b) for the flat-top seed of duration $6,000\omega_0^{-1}$ with $\lambda_1 = 1.638\lambda_0$.




Figure 2.21: The time-average measured gain seen in simulations as we vary the continuous seed intensity using seeds with $\lambda_1 = 1.644\lambda_0$ (dashed blue) and $1.658\lambda_0$ (solid red). For comparison, we use horizontal dashes on the left side of the plot to mark the steady-state gain from linear relativistic PIC theory.

seeds at different wavelengths and intensities, and examining the results using plots of the reflected light.

Figure 2.21 shows the time-average measured gain for various seed intensities with $\lambda_1 = 1.644\lambda_0$ and $1.658\lambda_0$. When we use a seed with $\lambda_1 = 1.644\lambda_0$, an intensity of $1.25 \times 10^{-4}I_0$ (smallest value shown) is enough to cause inflation. However, when we use a seed with $\lambda_1 = 1.658\lambda_0$, we do not see inflation until the seed intensity reaches $8 \times 10^{-3}I_0$. The measured gain at both wavelengths decreases with seed intensity due to pump depletion once inflation sets in.

In order to understand the results with a $1.658\lambda_0$ seed, we first discuss trapping

effects on the SRS gain spectrum via Equation 2.10. Trapping nonlinearity leads to a reduction in χ_i and thus Landau damping, as well as a down-shift in the natural plasma wave frequency. This latter effect decreases the resonant scattered wavelength, where $1 + \chi_r = 0$. In electrostatic simulations with an external driver, Fahlen showed behavior consistent with this picture [Fah10]. For $k_2\lambda_{De} \sim 0.3$ and a driver frequency above the natural frequency, the plasma response is even smaller than the linear response. However, when the driver is below the natural frequency, a larger response is obtained.

Trapping, therefore, makes the gain spectrum narrower and peaked at a smaller wavelength. The gain increases for wavelengths near resonance. However, for wavelengths far from resonance, the gain does *not* increase, and, in fact, it scales as χ_i , so that it vanishes as $\chi_i \rightarrow 0$. We can see this effect clearly by examining Equation 2.10. Since the seed is far from resonance, $1 + \chi_r \gg \chi_i$, so $g_0 \propto \chi_i / (1 + \chi_r)^2$. This vanishing gain is clear in Figure 2.22a, where the gain approaches zero in steady state for the continuous $1.658\lambda_0$ seed with $I_{1s}/I_0 = 4 \times 10^{-3}$. Thus, inflation is not possible at non-resonant wavelengths. For hot, low-density plasmas, there is no resonant wavelength, and all phase-matched plasma waves satisfy the loss of resonance condition $k_2\lambda_{De} > 0.53$ [RR01]. Inflation cannot occur in such a plasma at any scattered wavelength.

The red curve in Figure 2.21 is for a seed wavelength that is non-resonant and larger than the linear resonance. The reduction of χ_i due to trapping cannot lead to inflation at this wavelength, and the nonlinear frequency shift will move the resonance farther away. Both effects conspire to reduce the SRS gain below its linear value, which is what we observe for the lowest seed intensities. The number of bounce orbits completed by resonant electrons [SWL07, SWR12], based on the plasma wave amplitude computed from linear theory, is > 4 . It is thus consistent for trapping nonlinearity to occur and reduce the SRS gain. The increase in gain for $I_{1s}/I_0 = 8 \times 10^{-3}$, as shown in Figure 2.22b, first develops at linear resonance,

$1.644\lambda_0$, *not* at the seed value of $1.658\lambda_0$, then shifts down in wavelength with time to finish around $1.638\lambda_0$. This progression is similar to the one in Figure 2.16. We performed a similar run using a flat-top seed of duration $1,000\omega_0^{-1}$ with a central wavelength of $1.658\lambda_0$ and $I_{1s}/I_0 = 4 \times 10^{-3}$ and observed inflation similar to that in Figure 2.22b, but without the oscillations, which raises the possibility that continuous seeds can suppress inflation.

There are oscillations in Figure 2.22b, which occur because the seed is driving SRS off resonance. As we discussed earlier, this non-resonant drive leads to oscillations in the reflected light with a period of $2\pi/\Delta\omega_{NR}$, where $\Delta\omega_{NR}$ is the difference between the seed frequency and the resonant frequency, at which inflationary SRS grows. Equivalently, the $\Delta\omega_{NR}$ is the difference between the seeded beat drive frequency and the frequency of the plasma wave packet. When we use a seed with $I_{1s}/I_0 = 1.024$, the oscillations are more prominent and faster, and we see no amplification, as seen in Figure 2.22c. The increased oscillation frequency is due to the higher amplitude plasma wave undergoing a greater frequency shift. We also see a beat wave pattern covering many oscillations, which is caused by the nonlinear frequency shift of the plasma wave packet, as described earlier in this section.

Figure 2.23 shows the reflected light in simulations using a (low-, moderate-, high-) intensity seed with $I_{1s}/I_0 = (1.25 \times 10^{-4}, 8 \times 10^{-3}, 1.024)$, but $\lambda_1 = 1.644\lambda_0$. When we use the low-intensity seed, we see the reflected light increase monotonically until it saturates near the end of the simulation. Unlike when we used $\lambda_1 = 1.658\lambda_0$, this seed is near resonance, so $1 + \chi_r \ll \chi_i$, and $g_0 \propto 1/\chi_i$. Therefore, particle trapping increases the gain, as explained earlier in the Chapter. Oscillations begin to appear again in the simulation with the moderate-intensity seed, which indicates that the seed is slightly off resonance, as we expect to occur as the plasma wave undergoes a nonlinear frequency shift. These oscillations once again become more prominent when we use the high-intensity seed, and we see a




Figure 2.22: The reflected light in simulations using a continuous seed with $\lambda_1 = 1.658\lambda_0$ (solid red in Figure 2.21). $I_{1s}/I_0 = 4 \times 10^{-3}$ (a), 8×10^{-3} (b), and $1.024I_0$ (c). The horizontal dash-dotted line indicates the amplitude of the seed (unamplified) reflected light.




Figure 2.23: The reflected light in simulations using a continuous seed with $\lambda_1 = 1.644\lambda_0$ (dashed blue in Figure 2.21). $I_{1s}/I_0 = 1.25 \times 10^{-4}$ (a), 8×10^{-3} (b), and 1.024 (c). The horizontal dash-dotted line indicates the amplitude of the seed (unamplified) reflected light.

beat wave pattern appear again.

2.9 Conclusions and Future Research

Using 1D OSIRIS PIC simulations, we have studied BSRS of a well-defined seed pulse with variable shape, intensity, and wavelength. We found that backward Raman amplification of a seed scattered light pulse can remain in the strongly damped convective regime. That is, for a sufficiently weak seed, kinetic inflation does not occur. Peak seed amplification occurs near the peak of the linear gain spectrum when we take into account special relativity and PIC effects such

as finite particle size, finite-difference operators, and field smoothing. Detailed comparisons with linear coupled-mode predictions for envelope dynamics show excellent agreement. If the seed pulse is intense enough, the driven plasma wave traps particles, thereby lowering the Landau damping after a time on the order of a bounce period. The plasma wave continues to scatter light and grows after the seed leaves the system, leading to kinetic inflation later in the simulation. When we extend the seed in time, kinetic inflation occurs while the seed is still present, and we measure dynamic seed amplification, which significantly exceeds the linear gain rate at times, and can also turn negative at other times (i.e., it transfers energy into the pump). When we use a continuous seed, we find that the onset of inflation depends strongly on the seed wavelength. When the continuous seed wavelength is near the peak of the gain curve, we see kinetic inflation occur with very low seed amplitudes, while higher seed amplitudes are necessary to drive kinetic inflation using a non-resonant seed.

Connecting kinetic inflation with experiments involving lasers with many hot spots or speckles is still at an early stage. Past research has demonstrated that hot electrons, beam acoustic modes, and side-scattered light can couple hot-spots transversely in 2D simulations, leading to higher BSRS from each hot spot than one would see without coupling [BWL06, YAR12a]. Our results suggest that we can expect to see inflationary scattering from the lower-intensity parts of the beam interacting with flattened (non-Maxwellian) distributions or plasma waves generated in more intense parts, or increased scattered light from a few speckles triggering a “chain reaction” of downstream inflation. If this chain reaction occurs, most of the Raman in underdense laser-produced plasmas, such as ICF targets, will be inflationary. PIC codes can be used to model hundreds of speckles, but to simulate large volumes across a hohlraum effectively requires the use of envelope codes (such as pF3D), for which reduced models of kinetic nonlinearity are being pursued by several groups [BSG09, YF09, LCW07].

An important factor we have not explored is de-trapping mechanisms. All of our simulations in this Chapter are 1D and “collisionless,” which means that plasma waves easily trap electrons, and trapped electrons cannot leave sideways or be scattered out of the plasma wave. Since particle trapping is necessary for inflationary scattering, any de-trapping mechanisms make it less likely to occur and impose a threshold amplitude for inflation (the threshold in the present work is set by the finite amplitude and duration of the seed). In particular, collisions can kick electrons out of a plasma wave’s potential well, and electrons can traverse the plasma wave in less than one bounce period in higher dimensions [FWG11, SWR12]. Future research in this area should explore the effect of these de-trapping processes on inflation, and validate reduced descriptions.

CHAPTER 3

Studies of Particle Wake Potentials in Plasmas

3.1 Introduction

The stopping power of an energetic electron in a plasma is the result of interactions (collisions) between the moving electron and background plasma electrons. These interactions can also lead to the excitation of coherent (Čerenkov) wakes. Collisions affect the stopping power in a plasma when particles are closer than distances smaller than v/ω_{pe} where v is the test particles speed and ω_{pe} is the plasma frequency. To accurately model the collisions, interactions down to distances smaller than the classical distance of closest approach or the de Broglie wavelength (which ever is larger) need to be resolved. In the standard PIC method [Daw83, BL85], the cell (particle) size is on the order of the Debye length. Therefore, the standard PIC method cannot be used to model collisions from first principles. The particle-particle particle-mesh (PPPM) method [HGE73] is useful for Coulomb (electrostatic) interactions. It models long range interactions using the standard PIC method and interactions within the Debye length by summing up the Coulomb forces from each particle for close encounters. If one simulation electron is used per real electron, then the PPPM method can be used to study Coulomb collisions from first principles. If interactions at scales less than the de Broglie wavelength are important, then quantum mechanics needs to be included.

In the absence of quantum effects, the PIC method can also be used to study collisions from first principles if the cell (particle) size is made sufficiently small

and one simulation electron is used for each real electron. The cell (particle) size can be gradually reduced to see if the results converge. While the PPPM method may be more efficient than the PIC method (when the interactions with very small impact parameters are important) for studying Coulomb collisions, it has limitations for studying the stopping of relativistic electrons.

When an electron moves relativistically, its own self-fields are no longer described by electrostatics, but are electromagnetic in character. However, if it is moving towards a stationary electron it feels a purely electrostatic force (if the impact parameter is large enough then the electron remains nearly stationary). For the general case, where the background electrons are not stationary (the plasma is hot) and where small impact parameters matter (so the background electrons are significantly deflected), then one cannot use electrostatic fields to model the interactions between these particles. In such cases, the PPPM method cannot be used in its current electrostatic form. For relativistic interactions and small impact parameters, quantum electrodynamic (QED) effects need to be included.

Later in this dissertation, we will investigate the use of electromagnetic (full or what we call quasi-static) PIC codes with small cell (particle) sizes to study the stopping of relativistic electrons in PIC codes. In this regime, the wakefield generated by the relativistic electron is very important, and the wakefields of two or more electrons can potentially interact coherently to greatly increase the stopping power of a “beam” of relativistic electrons.

To set the stage for these studies, in this Chapter we will compare results from standard PIC and PIC with small cell sizes with those from a PPPM code for purely electrostatic interactions. We will examine the details of the wake using the subtraction technique and compare the simulation results against theory.

3.2 PIC & PPPM Methods

The PIC method [Daw83, BL85] is simple in concept. The point of the method is to accurately model collective behavior of plasmas while achieving a speed-up over the traditional molecular dynamics algorithm by smoothing out interactions that occur inside a grid spacing (often a Debye length). We achieve this smoothing and speed-up by solving for the fields on a spatial grid rather than solving for the force on each particle from all the other particles directly. Since PIC uses a grid, we often say that the particles have a finite size. The particle size and the grid size can be viewed separately, which leads to the view that the mathematical model for PIC is the interaction of finite-size particles. Typically, the grid spacing is on the order of a Debye length, so interactions on smaller spatial scales are unresolved. However, the grid spacing can be smaller [DMD94].

Each iteration in the PIC method contains four steps. First, given the positions and velocities of the particles in the system, interpolate the charge and current to the grid points. Second, solve for the electric and magnetic fields on the grid. Third, interpolate to find the values of the electric and magnetic fields at the particle positions. Finally, push the particles using the Lorentz force law. The process is demonstrated visually in Figure 3.1. This process gives us an $O(N)$ algorithm, where N is the number of particles, as opposed to the $O(N^2)$ algorithm of molecular dynamics.

The PPPM method [HGE73] is designed to maintain computational efficiency and take into account pair-wise interactions between particles. It calculates the smooth long-range forces using a grid and the PIC method for computational efficiency. It then computes pair-wise interactions within a cut-off sphere specified by the user. Computing the pair-wise interactions typically makes the PPPM method more computationally intensive than “standard” PIC. The pair-wise forces from the PIC method need to be subtracted first. We have again demonstrated




Figure 3.1: The cycles used in the PIC the PPPM methods. Conceptually, the two cycles are identical except for accounting for pair-wise interactions (red box) in the PPPM method.

the steps visually in Figure 3.1. The PPPM method has the advantage over PIC of more accurately simulating processes in which Coulomb collisions play an important role. However, as noted earlier, the PPPM method has not yet been extended to fully electromagnetic problems.

Collisions in plasmas are commonly described using a Fokker-Planck operator, which implicitly assumes that large-angle scattering events are not important. The velocity drag and diffusion coefficients are then calculated using the Landau-Boltzmann or Leanard-Balescu approaches. The relationship between these two approaches is well understood and it can be shown that the Landau-Boltzmann can be recovered from the Leanard-Balescu under certain assumptions. In the Leanard-Balescu approach, the “wakes” made by the test charge are central. Therefore, it is important to perform a detailed study of the wakes made by test charges for the PIC and PPPM methods.

3.3 Collisionless Kinetic Theory of Particle Wakes

We derive the time-dependent potential of a particle in a plasma following Decyk [Dec87]. The potential generated by a charge density ρ in a plasma (in CGS units) is

$$\phi(\vec{k}, \omega) = \frac{4\pi\rho(\vec{k}, \omega)}{k^2\varepsilon(\vec{k}, \omega)}, \quad (3.1)$$

where $\varepsilon(\vec{k}, \omega)$ is the collisionless plasma dielectric from Vlasov (henceforth collisionless) kinetic theory. If we assume an electron plasma with a Maxwellian velocity distribution and immobile ions, the dielectric is given by

$$\varepsilon(\vec{k}, \omega) = 1 - \frac{\omega_{pe}^2}{2k^2v_{th}^2} Z' \left(\frac{\omega}{\sqrt{2}kv_{th}} \right), \quad (3.2)$$

where $v_{th} = \sqrt{T_e/m_e}$ is the electron thermal speed and $Z'(s) = dZ/ds$. $Z'(s)$ must be calculated numerically, and is typically found by first computing $Z(s)$, which is known as the plasma dispersion function. The two are related by the differential equation, $Z'(s) + 2sZ(s) + 2 = 0$. $Z(s)$ is related to the error function by

$$Z(s) = i\sqrt{\pi}e^{-s^2}[1 + \operatorname{erf}(is)]. \quad (3.3)$$

Suppose we insert a test particle with charge q_t at \vec{x}_0 with velocity \vec{v} at $t = 0$. Then, the charge density takes the form,

$$\rho(\vec{x}, t) = q_t\delta(\vec{x} - \vec{x}_0 - \vec{v}t), \quad t > 0 \quad (3.4)$$

in real space and, performing a Fourier transform in position and a Laplace transform in time, we obtain the test charge in transform space,

$$\rho(\vec{k}, \omega) = \frac{q_t e^{-i\vec{k}\cdot\vec{x}_0}}{i(\vec{k}\cdot\vec{v} - \omega)}. \quad (3.5)$$

Therefore,

$$\phi(\vec{k}, \omega) = \frac{4\pi q_t e^{-i\vec{k}\cdot\vec{x}_0}}{ik^2(\vec{k}\cdot\vec{v} - \omega)\varepsilon(\vec{k}, \omega)}. \quad (3.6)$$

We need $\phi(\vec{x}, t)$, so we first invert the Laplace transform, which is defined as

$$\phi(\vec{k}, t) = \frac{4\pi q_t e^{-i\vec{k}\cdot\vec{x}_0}}{k^2} \int_{-\infty-ic}^{\infty-ic} \frac{d\omega e^{-i\omega t}}{2\pi i(\vec{k}\cdot\vec{v} - \omega)\varepsilon(\vec{k}, \omega)}, \quad t > 0. \quad (3.7)$$

The integrand has poles at $\omega = \vec{k}\cdot\vec{v}$ and $\omega = \omega_j(\vec{k})$, where $\omega_j(\vec{k})$ are the roots of $\varepsilon(\vec{k}, \omega) = 0$. Therefore, we can apply Cauchy's residue theorem, which yields,

$$\phi(\vec{k}, t) = \frac{4\pi q_t e^{-i\vec{k}\cdot(\vec{x}_0 + \vec{v}t)}}{k^2 \varepsilon(\vec{k}, \vec{k}\cdot\vec{v})} - \sum_j \frac{4\pi q_t e^{-i\vec{k}\cdot\vec{x}_0} e^{-i\omega_j t}}{k^2 (\vec{k}\cdot\vec{v} - \omega_j) \frac{\partial \varepsilon(\vec{k}, \omega)}{\partial \omega} \Big|_{\omega=\omega_j}}, \quad t > 0. \quad (3.8)$$

We next apply the inverse Fourier transform and write $\phi(\vec{x}, t) = \phi_D(\vec{x}, t) + \phi_C(\vec{x}, t)$, where

$$\phi_D(\vec{x}, t) = 4\pi q_t \int_{-\infty}^{\infty} \frac{d\vec{k}}{(2\pi)^n} \frac{e^{i\vec{k}\cdot(\vec{x} - \vec{x}_0 - \vec{v}t)}}{k^2 \varepsilon(\vec{k}, \vec{k}\cdot\vec{v})}, \quad t > 0 \quad (3.9)$$

and

$$\phi_C(\vec{x}, t) = -4\pi q_t \sum_j \int_{-\infty}^{\infty} \frac{d\vec{k}}{(2\pi)^n} \frac{e^{i\vec{k}\cdot(\vec{x} - \vec{x}_0)} e^{-i\omega_j t}}{k^2 (\vec{k}\cdot\vec{v} - \omega_j) \frac{\partial \varepsilon(\vec{k}, \omega)}{\partial \omega} \Big|_{\omega=\omega_j}}, \quad t > 0, \quad (3.10)$$

where n is the spatial dimensionality. Decyk calls $\phi_D(\vec{x}, t)$ the Debye cloud term and $\phi_C(\vec{x}, t)$ the Čerenkov term. The ϕ_D term is the wake while the ϕ_C term ensures the correct initial and boundary conditions. We shall see that we usually only need to keep the least damped root of $\varepsilon(\vec{k}, \omega)$, even though $\phi(\vec{x}, t)$ contains a sum over all roots.

3.4 The Subtraction Technique

Fluctuations occur in real and simulated plasmas due to discreteness effects and are not part of the smooth kinetic theory description. If the temperature of the plasma is high enough, fluctuations can be large compared to the wake of a single particle. Particle simulations such as PIC and PPPM include fluctuations, so we must remove them from the data during analysis. In PIC simulations, despite the use of finite-size particles, these fluctuations can be enhanced when a small number of particles are used. The subtraction technique [Dec87] removes the




Figure 3.2: The electrostatic potential (units e/λ_{De}) from the BEPS simulation of Section 3.5, with a test charge and an additional neutralizing charge (left), one without the test charge and its neutralizing charge but otherwise identical (center), and the difference of the two (right). The test charge is at $\vec{x} = (128\lambda_{De}, 128\lambda_{De}, 214\lambda_{De})$ at the time shown.

uncorrelated background through the use of two simulations. The first simulation includes the small perturbation, such as a test charge, while the second does not. The simulations are otherwise identical. Subtracting the second from the first reveals the response to the perturbation, as shown in Figure 3.2. Decyk used this method to show that, even when a few simulation particles are used, the wakes initially created by charges in PIC codes agree with the collisionless kinetic theory calculations.

3.5 PIC Results and Comparison with Collisionless Kinetic Theory

We first study particle wakes by performing traditional PIC simulations using the 3D periodic electrostatic code BEPS [DN04]. BEPS solves for the potential and the electric fields in the plasma using FFTs, supports linear and quadratic particle shapes, and offers a relativistic and non-relativistic particle pusher. We use quadratic particle shapes and a non-relativistic particle pusher in all BEPS simulations used in this Chapter.

BEPS uses a normalized set of units. Time is normalized to the plasma frequency and distance is normalized to the cell width, which can be easily related to the Debye length. Charge is normalized to the absolute value of the charge on an electron and mass is normalized to the mass of an electron. Each simulation electron has the correct charge to mass ratio, e/m_e , but it typically has a charge and mass that are much larger than a real electron. Therefore, each simulation corresponds to a family of cases (different densities) in which collisionless phenomena that only depend on q/m , qn_0 , and mn_0 , are important. However, collisional or discrete particle effects that depend on q^2/m or $n_0\lambda_D^3$ will vary within the family of simulations. Since we are interested in both collisional and collisionless effects, we give $n_e\lambda_{De}^3$ and Δ/λ_{De} for each simulation.

Figure 3.3 shows the results of the BEPS simulation. The box has dimensions of $256 \times 256 \times 512$ cells, where the width of a cell and the particle size are a Debye length. We use 64 negatively-charged background particles per cell ($n_e\lambda_{De}^3 = 64$), each carrying an electron charge, which are neutralized by a smooth positive background. The test particle also carries an electron charge and starts at $\vec{x}_0 = (128\lambda_{De}, 128\lambda_{De}, 64\lambda_{De})$ with *fixed* $\vec{v} = 3v_{th}\hat{z}$. A neutralizing charge is placed at $\vec{x}_1 = (254\lambda_{De}, 254\lambda_{De}, 510\lambda_{De})$ and held fixed. We can see the contribution of the Čerenkov term centralized around \vec{x}_0 , similar to the ripples on the surface of

a pond spreading out from the point where a pebble enters the water. The waves immediately following the test charge as it moves across the box come from the Debye term. We can see the effect of the Debye term better in Figure 3.4, in which we vary the speed of the test charge while keeping all other parameters the same.

We calculate the theoretical potentials, ϕ_C and ϕ_D , on a mesh with the same dimensions as the box. The test charge at its initial position, the test charge at the current position, and the neutralizing charge are deposited to separate meshes using second-order B-splines, just as charges are deposited to the mesh in BEPS. We performed a Fast Fourier Transform on the deposited charge, which we used in Equation 3.8. The roots of the plasma dielectric (3.2) were found using the continuous fraction method of Derfler and Simonen [DS69]. As mentioned earlier, we found that we only needed to keep the least damped root. Finally, we transformed the potential to real space. Figure 3.5 shows the result along with the result from the simulation.

The direct comparison of the simulation results with theory in Figure 3.5 yields decent results, but we can do better if we take into account two known aspects of PIC codes. The first effect is the finite size of the plasma particles.

$$q\delta(\vec{x}) \rightarrow qS(\vec{x}), \quad (3.11)$$

where $S(\vec{x})$ is known as the “shape factor.” The shape factor is also viewed as the interpolating function, and needs to be included in the plasma dispersion relation [BL85]. Transforming to Fourier space,

$$qS(\vec{x}) \rightarrow qS(\vec{k}), \quad (3.12)$$

where, for second order B-splines,

$$S(\vec{k}) = \prod_i \left(\frac{\sin(k_i \Delta/2)}{k_i \Delta/2} \right)^3, \quad (3.13)$$




Figure 3.3: The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in the BEPS simulation of Section 3.5. The z position of the test charge is $91\lambda_{De}$ and $334\lambda_{De}$ in the left and right plots, respectively.




Figure 3.4: The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in the BEPS simulation of Section 3.5 with various test particle speeds. The test particle speed is $3v_{th}$ and $10v_{th}$ and its z position is $154\lambda_{De}$ and $364\lambda_{De}$ in the left and right plots, respectively.




Figure 3.5: The electrostatic potential along the trajectory of the test charge with speed $3v_{th}$ in the BEPS simulation of Section 3.5 (solid) and as predicted by collisionless kinetic theory (dotted) at $20\omega_{pe}^{-1}$. The vertical line indicates the position of the test charge.

where Δ is the width of a cell and $i = x, y, z$. Second, the force on a particle in a PIC code depends on both its distance from other particles and the grid. The dependence of the force on a particle's distance from the grid can lead to a self-heating instability when the cell size is too large. PIC codes which solve for the fields in Fourier space, such as BEPS, can mitigate this effect by filtering the charge density and the force in Fourier space. We need to apply this filter function after transforming the charge to Fourier space in the theoretical calculation. BEPS uses a Gaussian filter,

$$F(\vec{k}) = e^{-(k\Delta a)^2/2}, \quad (3.14)$$

where $a = 0.866667$ in our simulations. The filter and shape functions only affect the plasma frequency in the dispersion relation, so that

$$\omega_{pe}^2 \rightarrow \omega_{pe}^2 S^2(\vec{k}) F^2(\vec{k}), \quad (3.15)$$

where $S(\vec{k})$ and $F(\vec{k})$ are squared because they are applied to both the charge density and the force. We note that $S(\vec{k})F(\vec{k})$ can be viewed as an overall shape function. After taking these two factors into account in the theoretical calculation, the agreement with simulation becomes much better, as seen in Figure 3.6. At late times, we start seeing ripples in the simulation result, which arise due to discreteness effects.

3.6 PPPM Results and Comparison with Traditional PIC

The PPPM code ddcMD [GGB12, RGC09], like BEPS, is 3D, periodic, electrostatic, and non-relativistic. The code offers the choice of several different pseudo-potentials, but we use a pure Coulomb potential in these simulations.

The PPPM method is much more computationally expensive than traditional PIC because it requires much shorter time-steps because it also includes N^2 calculations to accurately resolve the pair-wise interactions that are unresolved in




Figure 3.6: The electrostatic potential along the trajectory of the test particle with speed $3v_{th}$ at various times in the BEPS simulation of Section 3.5 (solid) and as predicted by collisionless kinetic theory taking into account finite particle size and Fourier space smoothing (dash-dotted). The vertical lines indicate the position of the test charge.

traditional PIC. Therefore, the ddcMD simulations are much smaller and much shorter than the BEPS simulation in the previous section. The box is $3.85\lambda_{De}$ on each side and $n_e\lambda_{De}^3 = 36,700$, so the plasma is weakly coupled, and we should expect PIC simulations to accurately describe the wake even for cell sizes $\sim \lambda_{De}$. As in the BEPS simulation, the background electrons are neutralized by a smooth positive background. The test particle is an electron starting at $\vec{x}_0 = (1.925\lambda_{De}, 1.925\lambda_{De}, 0.105\lambda_{De})$ and again moving in the \hat{z} direction with a fixed speed of $3v_{th}$, but we do not include a neutralizing charge this time. In Section 3.5, the neutralizing charge served to prevent an overall difference in potential between the two simulations involved in the subtraction technique. However, by subtracting the free-space potential of the test charge in this Section, as discussed below, we account for this difference. We choose to use 250,000 time-steps per ω_{pe}^{-1} to make sure we are accurately resolving collisions with small impact parameters, $b \ll \lambda_{De}$.

Figure 3.7 shows the evolution of the projectile wake with time. In addition to subtracting out the background noise, we subtract the free-space potential of the test particle, which would otherwise overwhelm any interesting features we wish to observe on these small spatial and temporal scales. Notice what look like dipoles appearing at later times, which we shall suggestively call “collisional bubbles.” The collisional bubbles appear because, when two particles collide, a small change in their initial positions can result in a large change in their final positions, as illustrated in Figure 3.8. We are observing the effect of path differences using the subtraction technique, and collisions are simply amplifying these path differences. By $t=1.1\omega_{pe}^{-1}$, the path differences between the two runs are so great that we have difficulty seeing the wake.

We now compare the results of a BEPS simulation with the ddcMD simulation. The BEPS simulation has $128 \times 128 \times 128$ cells and one background particle per cell so that the number of particles in both simulations is the same. The charge of the background particle is neutralized by a smooth positive background as in




Figure 3.7: The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in a ddcMD simulation. The z position of the test charge is $0.105\lambda_{De}$, $0.705\lambda_{De}$, $1.785\lambda_{De}$, and $3.405\lambda_{De}$ in the top left, top right, bottom left, and bottom right plots, respectively.




Figure 3.8: Collisions involving two electrons each, illustrating that a small change of the electron starting positions leads to a large change in their positions at later times. The electrons in the collision shown by the solid lines start at $(-10, 0.1)$ and $(10, -0.1)$ while those in the collision shown by the dotted lines start at $(-10, 0.2)$ and $(10, -0.2)$.




Figure 3.9: The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in a BEPS simulation using the same parameters as the ddcMD simulation in Figure 3.7. The z position of the test charge is $0.705\lambda_{De}$ in the left plot and $3.405\lambda_{De}$ in the right plot.

the previous BEPS simulation. The box has dimensions of $3.85\lambda_{De}$ per side, like the ddcMD simulation, and the test charge is placed at the same initial position as in the ddcMD simulation, again with $\vec{v} = 3v_{th}\hat{z}$. The time-step is $0.05\omega_{pe}^{-1}$. The result is in Figure 3.9, where we can see that the wake evolves almost identically to the wake in the ddcMD simulation of Figure 3.7, except that the BEPS simulation has no collisional bubbles.

3.7 Fine-Mesh PIC

Given a fine enough mesh and a small enough time-step, PIC should be able to reproduce the results of PPPM simulations. The necessary mesh size depends on the importance of impact parameters smaller than the particle (cell) size. We begin with a basic attempt to resolve collisions in a BEPS simulation. Figure 3.10 shows the results of a BEPS simulation using 2048 cells and $10\lambda_{De}$ per side with $n_e\lambda_{De}^3 = 56,623$, yielding approximately 152 cells per particle. Test particle starts at $(5\lambda_{De}, 5\lambda_{De}, 2.5\lambda_{De})$ with $\vec{v} = 3v_{th}\hat{z}$. A neutralizing charge is placed at




Figure 3.10: The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in a BEPS simulation with approximately 152 cells per particle. The z position of the test charge is $3.94\lambda_{De}$ in the left plot and $7.42\lambda_{De}$ in the right plot.

$(10\lambda_{De}, 10\lambda_{De}, 10\lambda_{De})$ and held fixed. The time-step is $10^{-4}\omega_{pe}^{-1}$. Once again, we subtract the free-space potential of the test charge when analyzing the results. If we look closely, we can see very small bubbles in the plots, but nothing as large as we see in the ddcMD simulations. We might see better results with a finer mesh and smaller time-step. We also note that we did not use a diagnostic in ddcMD to verify that binary collisions rather than another effect was causing the collisional bubbles. Such a diagnostic would determine the validity of our hypothesis that a lack of resolution in our PIC simulations is causing the discrepancy between our PIC and ddcMD results.

Figure 3.11 shows results of a BEPS simulation with parameters identical to the ddcMD simulation of Section 3.6. We loaded the initial particle positions from the ddcMD simulation into BEPS in an attempt to replicate the results. We now use 1024 cells and $3.85\lambda_{De}$ per side with $n_e\lambda_{De}^3 = 36,700$, giving us 512 cells per particle, along with a time-step of $2 \times 10^{-5}\omega_{pe}^{-1}$, and the test charge once again starts at $(1.925\lambda_{De}, 1.925\lambda_{De}, 0.105\lambda_{De})$ with $\vec{v} = 3v_{th}\hat{z}$. Many more bubbles are




Figure 3.11: The evolution of the electrostatic potential (e/λ_{De}) in a plane containing the path of the test particle in a BEPS simulation with 512 cells per particle. The z position of the test charge is $0.705\lambda_{De}$ in the left plot and $3.405\lambda_{De}$ in the right plot.

now visible, but there are still none as large as we were seeing in the ddcMD simulation. We may need yet a finer mesh.

3.8 Conclusions

We have shown that traditional electrostatic PIC simulations produce single particle wakes that agree with the theoretical predictions of collisionless kinetic theory in three dimensions even when few particles are used. We have included finite size particle effects into the theory, including k-space smoothing. The simulation wakes are determined using the subtraction technique. We also found that the details of the subtracted wake from PPPM and PIC with fine meshes (smaller particles) differ from those from traditional PIC simulations. The subtracted wake includes small bubbles which look like tiny dipoles, which are believed to arise when the trajectory of plasma particles change due to a close encounter. These differences do not necessarily affect the local electric field at the test par-

ticle, which is responsible for the stopping power; and this is studied in the next Chapter.

Directions for future work include understanding in detail the trajectories of particles for traditional PIC, fine cell PIC, and PPPM. It will also be interesting to compare the wakes predicted by electrostatic interactions with those to be described later for relativistic electrons including the full Maxwell equations. An intermediate step would be to study the wakes using a Darwin field solver [KR71, NL76, BL85]. This includes the magnetic fields of moving charges as well. The Darwin fields do not include retarded time effects, so a PPPM type model could be developed for the Darwin fields. Although we study the stopping power of electrons in later Chapters, much can be learned by studying the single particle wakes for different resolution using the PIC codes.

CHAPTER 4

Single-Particle Stopping Power

4.1 Introduction

There has been much discussion about the stopping power formulas for relativistic electrons over the past several decades. The commonly accepted formula for relativistic electron stopping power is one based on quantum electrodynamics (QED). In this chapter, we cover several of the well-known formulas for both ion and electron stopping power for fast particles, discuss how these are altered by relativistic effects, and summarize how QED effects are included. We compare the predictions of these expressions against simulations, which include some or all of the physics in some cases.

4.2 Stopping Power Concepts

Stopping power is an N-body problem. A particle (test charge projectile) moving through a plasma slows down by interacting with many particles in the plasma. However, we do not have any analytic methods to directly calculate all these interactions, so we simplify the problem by breaking it into two pieces. The two contributing factors to the stopping power are the wake or plasma wave the charge produces as it moves through the plasma, $(d\mathcal{E}/ds)_{p.w.}$, and discrete binary collisions as it encounters particles with smaller impact parameters, $(d\mathcal{E}/ds)_c$.

The test charge is stopped through an axial electric field at the location of the

particle. One can calculate an electric field by using a moving point particle as a source in the Vlasov-Poisson system of equations. The solution to the Vlasov-Poisson system for a moving test charge with a speed, v_t , much larger than v_{th} is a plasma wave wake with a phase velocity, v_t . The wake is very weakly damped and its amplitude on-axis diverges if the integral in k -space (see Equation 3.9) is not cut off at large k . The cutoff at large k corresponds to limiting interactions at very short distances or impact parameters. At such distances, the Vlasov (“collisionless”) description is not appropriate and discreteness effects need to be included. These effects are generally handled realizing that very small impact parameters are not consistent with stopping power that is a differential, $d\mathcal{E}/ds$, since a particle can be stopped in one collision or it can be limited from quantum considerations. Therefore, $(d\mathcal{E}/ds)_{p.w.}$ comes from evaluating the electric field from the wake at the location of the particle where the k -space integral is limited to k_{max} .

Collisions contribute to the stopping power via the continuous exchange of energy in binary interactions as the particle passes near background particles. Therefore, $(d\mathcal{E}/ds)_c$ is evaluated from binary collisions of impact parameters between $b_{max} \sim 1/k_{max}$ and b_{min} . To derive the correct answer, we assume that both the wake and collisions contribute to the stopping power, and we add the two together,

$$\frac{d\mathcal{E}}{ds} = \left(\frac{d\mathcal{E}}{ds}\right)_{p.w.} + \left(\frac{d\mathcal{E}}{ds}\right)_c. \quad (4.1)$$

We note that the classical $(d\mathcal{E}/ds)_{p.w.}$ needs to also be the same as form as $(d\mathcal{E}/ds)_c$ for impact parameters between a cutoff b_c and b_{max} , where b_c corresponds to the range of the dressed potential of the moving test charge.

Conceptually, people sometimes find use in assuming that either the wake or collisions alone contribute to the stopping power of energetic particles, and introduce a minimum or maximum distance cut-off to produce a finite result. When examining the wake contribution alone, we need a minimum distance cut-off. Dif-

fering arguments have been made regarding the logarithmic divergence of the wake on the axis, both at the location of the test particle and behind it. The value of the wake at the test particle gives rise to the stopping power. Various arguments for what limits the unphysical divergence on-axis have been proposed. Some are based on classical physics, such as interparticle spacing and breakdown of linear theory, while others are based on quantum physics, such as the uncertainty principle. The collisional component is necessary to account for discrete interactions with the background plasma particles at small distances, where the Vlasov wake is questionable. The wake is the result of collective oscillations set up from the deflection of plasma electrons by the test charge (collisions). These small-scale interactions often require accounting for quantum mechanics, which lead to the Bethe and QED formulas for stopping power. This collisional description does not take into account the screening of intermediate particles, leading to the divergence when we use it to account for collisions with plasma particles outside a screening distance. In most derivations, the collisional and wake descriptions combine in an elegant fashion, and eliminate the divergences, because both have logarithmic divergences. The wake description accounts for the screening at large distances, and the collisional description accounts for the discrete particle physics missing from the wake description.

The question inevitably arises about the distance from the test charge at which the transition between where the particle wake or the collisional stopping dominates. We will see in Section 4.5 that we do not need to determine a precise transition. We assume the transition exists, and it usually disappears when we add the results of the two derivations together, because they both have logarithmic divergences. We note that another way to study the relative contribution between the Vlasov wake and collisions is to look behind the particle and see how energy is left in the plasma. If the energy is in the form of collective or average drifts, then the wake dominates. Or, if the energy is in the random thermal motion, then

collisions have dominated.

The final stopping power formula for a test charge projectile q_t in an electron plasma typically takes the form

$$\frac{d\mathcal{E}}{ds} = -\frac{\omega_{pe}^2 q_t^2}{v_t^2} \ln\left(\frac{b_{max}}{b_{min}}\right), \quad (4.2)$$

with additional corrections in some formulations, such as QED. It is also convenient at times to use the variable L_d , where the d means “drag,” such that

$$\frac{d\mathcal{E}}{ds} = -\frac{\omega_{pe}^2 q_t^2}{v_t^2} L_d. \quad (4.3)$$

The additional corrections can usually be included in L_d . The use of b in the logarithm comes from an integration over impact parameter, as will be demonstrated in Section 4.4. b_{max} and b_{min} are also associated with the screening distance and with the large scattering angle the cut-offs (i.e. distance of closest approach with a plasma particle).

Many textbooks [Che12, KT73] and the *NRL Plasma Formulary* [HNU11] simply note that b_{max} must be a screening distance and use the Debye length. The Debye length may suffice for an introduction to stopping power in plasmas. However, for energetic (non-thermal) test charges, the Debye length is not appropriate. As we saw in Chapter 3, the potential surrounding a charge moving through a plasma with $v > v_{th}$ is much more complex than the Debye potential. More detailed calculations, such as the one in Section 4.4, demonstrate that the dynamic screening length, $L_t = v_t/\omega_{pe}$, where v_t is the projectile speed, is a more accurate choice for b_{max} . We set $b_{max} = L_t$ in all our stopping power formulas.

In the purely classical stopping power calculation, b_{min} becomes the classical distance of closest approach,

$$d_{ca} = \frac{2q_t e(m_t + m_e)}{v_t^2 m_t m_e}, \quad (4.4)$$

where m_t is the mass of the test charge and m_e the mass of an electron. The final

formula is,

$$L_d = \ln \left(\frac{m_t m_e v_t^3}{2\omega_{pe} q_t e (m_t + m_e)} \right). \quad (4.5)$$

In the quantum case for non-relativistic energies, b_{min} is the de Broglie wavelength, $\lambda_{dB} = \hbar/2m_t v_t$. The stopping power is

$$L_d = \ln \left(\frac{2m_t v_t^2}{\hbar\omega_{pe}} \right). \quad (4.6)$$

This formula is known as the Bethe formula, and is a common basic formula for calculating quantum mechanical stopping power. In this case, some people like to say that plasmon emission is quantized since $\hbar\omega_{pe}$ appears in the formula.

We can get the first term of the relativistic QED electron stopping formula to be discussed in Section 4.6 by using the de Broglie wavelength in the center of mass frame as b_{min} . $\lambda_{dB}^* = \hbar/m_e c \sqrt{2(\gamma - 1)}$ [SB08], where γ applies to the test electron in the lab frame. Then,

$$L_d = \ln \left(\frac{\sqrt{2m_e v_t^2 \mathcal{E}}}{\hbar\omega_{pe}} \right), \quad (4.7)$$

with $\mathcal{E} = (\gamma - 1)m_e c^2$.

When particles are moving fast enough, $v_t \approx c$ for a large range of energies. Under this condition, we often focus our attention on a single plasma density, so L_d varies with energy while the coefficient in front of it is constant. Under these conditions, L_d is often called the “stopping number.”

4.3 Simple Relativistic Cold Fluid Wake Stopping Power

Useful formulas for the projectile wake and the associated stopping power of a relativistic test charge can be developed in real space using fluid theory. The formulas are limited because the wake calculation is linear, leading to the logarithmic divergence on-axis, as discussed in Section 4.2. However, if the test charge has a finite size, as in PIC codes, then the result on-axis has a logarithmic term involving the

size of the electron. Furthermore, the uncertainty principle gives the relativistic electron a “finite size.” However, the relationship between PIC finite-size particle potentials and quantum statistical potentials (which smooth out the Coulomb singularity) [JM07, GBB12] often used in MD simulations, is left for future research. We will see how the wake formula breaks down in Section 5.3. However, the wake formula provides a useful real-space conceptual picture along with scaling factors, and the stopping power formula demonstrates the relationship between the wake and stopping power. Furthermore, trailing electrons inside the wake can interact with the wake, showing the interactions are not simply binary.

We ultimately want the electric field of the projectile wake in the direction of motion, which we can use to calculate the stopping power directly. To calculate the wake, we follow the derivation of Katsouleas in Reference [KWC87] and work in cylindrical (r, z) coordinates.

Assume that the test projectile with charge q_t is moving through a plasma with velocity $\vec{v}_t = v_t \hat{z}$. Its charge density is

$$\rho(\vec{x}, t) = q_t \delta(\vec{x} - \vec{v}_t t) = q_t \delta(\vec{r}) \delta(z - v_t t), \quad (4.8)$$

where $\delta(\vec{r}) = \delta(r)/(2\pi r)$. We find the resulting plasma wave by using the linearized fluid equations and Maxwell’s equations,

$$\frac{d\vec{v}}{dt} = -e\vec{E}/m_e \quad (4.9)$$

$$\frac{\partial n_1}{\partial t} + n_0 \nabla \cdot \vec{v} = 0 \quad (4.10)$$

$$\nabla \cdot \vec{E} = -4\pi e n_1 + 4\pi \rho \quad (4.11)$$

$$\nabla \times \vec{E} = -\frac{1}{c} \frac{\partial \vec{B}}{\partial t} \quad (4.12)$$

$$\nabla \times \vec{B} = \frac{4\pi}{c} \vec{j} + \frac{1}{c} \frac{\partial \vec{E}}{\partial t}. \quad (4.13)$$

Taking the derivative with respect to time of Equation 4.10 and substituting in

Equation 4.9 yields

$$\frac{\partial^2 n_1}{\partial t^2} + n_0 \nabla \cdot (-e \vec{E}/m) = 0. \quad (4.14)$$

If we then replace $\nabla \cdot \vec{E}$ with Equation 4.11, we get the driven wave equation.

$$\frac{\partial^2 n_1}{\partial t^2} - en_0/m(-4\pi en_1 + 4\pi\rho) = 0. \quad (4.15)$$

Using Equation 4.8, the plasma frequency $\omega_{pe}^2 = 4\pi e^2 n_0/m_e$, and the fact that $\delta(z - v_t t) = \delta(t - z/v_t)/v_t$,

$$\frac{\partial^2 n_1}{\partial t^2} + \omega_{pe}^2 n_1 = \omega_{pe}^2 \frac{q_t}{e} \delta(\vec{r}) \delta\left(t - \frac{z}{v_t}\right) \frac{1}{v_t}. \quad (4.16)$$

This formula is just a simple harmonic oscillator whose solution is a Green's function,

$$n_1 = \omega_{pe} \frac{q_t}{e} \frac{\delta(\vec{r})}{v_t} \sin\left[\omega_{pe}\left(t - \frac{z}{v_t}\right)\right] \Theta\left(t - \frac{z}{v_t}\right), \quad (4.17)$$

where Θ is the unit step function. Here, we will add the condition that $\Theta(0) = 1/2$.

We now need to find \vec{E} using n_1 . Taking the time derivative of Equation 4.13 and plugging it into the curl of Equation 4.12 produces a wave equation relating \vec{E} to \vec{j} ,

$$\begin{aligned} \nabla \times \nabla \times \vec{E} &= -\frac{4\pi}{c^2} \frac{\partial \vec{j}}{\partial t} - \frac{1}{c^2} \frac{\partial^2 \vec{E}}{\partial t^2} \\ \nabla(\nabla \cdot \vec{E}) - \nabla^2 \vec{E} &= -\frac{4\pi}{c^2} \frac{\partial \vec{j}}{\partial t} - \frac{1}{c^2} \frac{\partial^2 \vec{E}}{\partial t^2} \\ \left(\frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \nabla^2\right) \vec{E} &= -\frac{4\pi}{c^2} \frac{\partial \vec{j}}{\partial t} - \nabla(\nabla \cdot \vec{E}). \end{aligned} \quad (4.18)$$

Combining Equation 4.9 and the fact that $\vec{j} = -n_0 e \vec{v} + \rho \vec{v}_t$,

$$\frac{\partial \vec{j}}{\partial t} = -n_0 e \frac{\partial \vec{v}}{\partial t} + \frac{\partial \rho \vec{v}_t}{\partial t} = \frac{e^2 n_0}{m_e} \vec{E} + \frac{\partial \rho \vec{v}_t}{\partial t}. \quad (4.19)$$

Substituting Equations 4.19 and 4.11 into Equation 4.18 produces

$$\left(\frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \nabla^2\right) \vec{E} = -\frac{4\pi e^2 n_0}{m_e c^2} \vec{E} - \frac{4\pi}{c^2} \frac{\partial \rho \vec{v}_t}{\partial t} - \nabla(-4\pi en_1 + 4\pi\rho). \quad (4.20)$$

We now must assume that $\vec{v}_t \rightarrow c\hat{z}$. Then, we can use the fact that $\frac{1}{c}\frac{\partial}{\partial t} + \frac{\partial}{\partial z} = 0$ to get,

$$\left(\frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \nabla^2 + \frac{\omega_{pe}^2}{c^2}\right)\vec{E} = \nabla(4\pi en_1). \quad (4.21)$$

Substituting in for n_1 using Equation 4.17,

$$\begin{aligned} &\left(\frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \nabla^2 + \frac{\omega_{pe}^2}{c^2}\right)\vec{E} \\ &= 4\pi q_t \omega_{pe} \nabla \left(\frac{\delta(\vec{r})}{c} \sin \left[\omega_{pe} \left(t - \frac{z}{c} \right) \right] \Theta \left(t - \frac{z}{c} \right) \right). \end{aligned} \quad (4.22)$$

Using the approximation that $\frac{1}{c^2}\frac{\partial^2}{\partial t^2} \approx \frac{\partial^2}{\partial z^2}$, breaking up the Laplacian as $\nabla^2 = \nabla_{\perp}^2 + \frac{\partial^2}{\partial z^2}$, and defining $k_{pe} \equiv \omega_{pe}/c$, we can then write Equation 4.22 as

$$(\nabla_{\perp}^2 - k_{pe}^2)\vec{E} = -4\pi q_t k_{pe} \nabla \left(\delta(\vec{r}) \sin \left[\omega_{pe} \left(t - \frac{z}{c} \right) \right] \Theta \left(t - \frac{z}{c} \right) \right). \quad (4.23)$$

Focusing on E_z , which is our primary interest, we can write,

$$(\nabla_{\perp}^2 - k_{pe}^2)E_z = 4\pi q_t k_{pe} \delta(\vec{r}) \frac{\partial}{\partial z} \left(\sin \left[\omega_{pe} \left(t - \frac{z}{c} \right) \right] \Theta \left(t - \frac{z}{c} \right) \right). \quad (4.24)$$

The solution for E_z can then be written as

$$E_z = G_R(\vec{r}) 4\pi q_t k_{pe}^2 \cos \left[\omega_{pe} \left(t - \frac{z}{c} \right) \right] \Theta \left(t - \frac{z}{c} \right). \quad (4.25)$$

$G_R(\vec{r})$ is the Green's function response to the Helmholtz equation,

$$(\nabla_{\perp}^2 - k_{pe}^2)G_R(\vec{r}) = \delta(\vec{r}), \quad (4.26)$$

with the solution

$$G_R(\vec{r}) = -\frac{1}{2\pi} K_0(k_{pe}|\vec{r}|), \quad (4.27)$$

where K_0 is the zeroth-order modified Bessel function of the second kind. Finally, we obtain the solution for E_z ,

$$E_z = -2q_t k_{pe}^2 K_0(k_{pe}|\vec{r}|) \cos \left[\omega_{pe} \left(t - \frac{z}{c} \right) \right] \Theta \left(t - \frac{z}{c} \right), \quad (4.28)$$

The wake therefore oscillates in $t - z/c$ and falls off quickly in r . However, the solution diverges as $r \rightarrow 0$, which is clearly unphysical.

The stopping power from this wake is simply the electric field (force) in the z direction from the wake evaluated at the particle,

$$\begin{aligned} \frac{d\mathcal{E}}{dz} = q_t E_z \left(r \rightarrow 0, t - \frac{z}{c} = 0 \right) &= -2q_t^2 k_{pe}^2 K_0(k_{pe}r \rightarrow 0) \frac{1}{2} \\ &\approx -q_t^2 k_{pe}^2 \ln \left(\frac{1}{k_{pe}r \rightarrow 0} \right) \\ &\approx -\frac{q_t^2 \omega_{pe}^2}{c^2} \ln \left(\frac{c}{\omega_{pe}r \rightarrow 0} \right). \end{aligned} \quad (4.29)$$

This expression, based solely on classical arguments, diverges. As noted earlier, this divergence can be reconciled from either the breakdown of linear fluid theory, using finite-size particles, as in PIC, or from quantum mechanical effects. In the case of finite-size particles, we can integrate the charge over the Green's function to get the resulting wakefield. The result will depend on the specific particle shape. However, since the particle size shows up inside the logarithm, we can simply modify this part of the expression. Furthermore, to set the stage for our use of PIC simulations, we use the cell size, Δ , for the particle size, leading to,

$$\frac{d\mathcal{E}}{dz} \approx -\frac{q_t^2 \omega_{pe}^2}{c^2} \ln \left(\frac{c}{\omega_{pe} \Delta} \right). \quad (4.30)$$

We will show in Section 4.7.3 that this formula agrees reasonably well with results from QuickPIC simulations. Furthermore, for the conditions we simulate, we can show that, as the particle size becomes significantly smaller than the interparticle spacing, that the result saturates as Δ is reduced. Specifically, when Δ becomes less than 1/4 of the interparticle spacing, the result stops diverging. This saturation indicates that linear fluid theory may be breaking down before QED effects need to be included.

4.4 Classical Bohr Stopping

The Bohr formula (which can be applied to ions and electrons) is one of the simplest classical stopping power formulas for energetic particles with non-relativistic

speeds. As we will see in Section 4.7.1, it works well in the classical regime for $v_t \gg v_{th}$. There are many ways of deriving it, and we will cover two here, starting with the derivation of Peter and Meyer-ter-Vehn [PV91].

As the charge slows down and $v_t \rightarrow v_{th}$, corrections start to appear. Many of these corrections can be accounted for by numerically performing the relevant integrals. Simplifications of the corrections have been published, one of which can be found in Reference [PV91]. The most accurate derivation of the corrections can be found in Reference [BPJ05]. However, none of the stopping power derivations account for the equilibration with the background plasma population, which occurs via stochastic heating and prevents the test charge from stopping completely.

4.4.1 Peter & Meyer-ter-Vehn Derivation

We begin with the wake formula from Equation 3.9,

$$\phi(\vec{x}, t) = 4\pi q_t \int_{-\infty}^{\infty} \frac{d\vec{k}}{(2\pi)^n} \frac{e^{i\vec{k}\cdot(\vec{x}-\vec{v}_t t)}}{k^2 \varepsilon(\vec{k}, \vec{k}\cdot\vec{v})}, \quad t > 0 \quad (4.31)$$

where ε is defined in Equation 3.2 and we have set $\vec{x}_0 = 0$. Then, if we assume that the test charge is moving in the \hat{x} direction, the stopping power is

$$-\frac{d\mathcal{E}}{dx} = -F_x = q_t \left. \frac{\partial \phi}{\partial x} \right|_{\vec{x}=\vec{v}_t t}. \quad (4.32)$$

Now, we define $X(s) = \Re Z'(s)$ and $Y(s) = \Im Z'(s)$. In this case, $s = \frac{\vec{k}\cdot\vec{v}_t}{\sqrt{2k}v_{th}} = \frac{v_t \cos(\theta)}{\sqrt{2}v_{th}}$, where θ is measured from the $+x$ -axis. Then,

$$L_d = 4\pi 2\pi \frac{v_t^2}{\omega_{pe}^2} \int_0^{\infty} \frac{k^2 dk}{(2\pi)^3} \int_{-1}^1 d \cos(\theta) \frac{ik \cos(\theta)}{k^2 \left(1 - \frac{\omega_{pe}^2}{2k^2 v_{th}^2} (X(s) + iY(s))\right)}. \quad (4.33)$$

We make a change of variables to simplify the calculation. Let $K = k\lambda_{De}$, $V_t =$

v_t/v_{th} , and $\mu = \cos(\theta)$. Now, $s = V_t\mu/\sqrt{2}$. Also using the fact that $v_{th}/\omega_{pe} = \lambda_{De}$,

$$L_d = 4\pi 2\pi V_t^2 \lambda_{de}^2 \int_0^\infty \frac{K^2 dK}{\lambda_{De}^3 (2\pi)^3} \int_{-V_t/\sqrt{2}}^{V_t/\sqrt{2}} ds \frac{\sqrt{2}}{V_t} \frac{i \left(\frac{K}{\lambda_{De}}\right) s \frac{\sqrt{2}}{V_t}}{\left(\frac{K}{\lambda_{De}}\right)^2 \left(1 - \frac{1}{2K^2} (X(s) + iY(s))\right)}. \quad (4.34)$$

Simplifying,

$$L_d = \frac{1}{\pi} \int_0^\infty K^3 dK \int_{-V_t/\sqrt{2}}^{V_t/\sqrt{2}} ds \frac{4is}{2K^2 \left(1 - \frac{1}{2K^2} (X(s) + iY(s))\right)}. \quad (4.35)$$

Distributing $2K^2$ in the denominator and multiplying the top and bottom by the complex conjugate of the denominator,

$$L_d = \frac{1}{\pi} \int_0^\infty K^3 dK \int_{-V_t/\sqrt{2}}^{V_t/\sqrt{2}} ds \frac{4is(2K^2 - (X(s) - iY(s)))}{(2K^2 - X(s))^2 + Y^2(s)}. \quad (4.36)$$

Given that $Z'(-s) = Z'^*(s)$, the imaginary part of the integrand is odd as a function of s , so it drops out. Therefore,

$$L_d = \frac{4}{\pi} \int_0^\infty K^3 dK \int_{-V_t/\sqrt{2}}^{V_t/\sqrt{2}} ds \frac{-sY(s)}{(2K^2 - X(s))^2 + Y^2(s)}. \quad (4.37)$$

We must cut off the K integral to avoid a logarithmic divergence. We set $K_{max} = \lambda_{De}/d_{ca}$, where d_{ca} is the classical distance of closest approach with a background electron. For an ion ($m_t \gg m_e$) and $V_t \gg 1$ (where the Bohr limit applies),

$$K_{max} = \frac{\lambda_{De}}{d_{ca}} = \frac{2\pi N_D V_t^2}{|Z_t|}, \quad (4.38)$$

where $Z_t = q_t/e$ and $N_D = n_0 \lambda_{De}^3$. For an electron projectile, d_{ca} doubles because we must use the reduced mass. Therefore,

$$L_d = -\frac{4}{\pi} \int_0^{K_{max}} K^3 dK \int_{-V_t/\sqrt{2}}^{V_t/\sqrt{2}} ds \frac{sY(s)}{(2K^2 - X(s))^2 + Y^2(s)}. \quad (4.39)$$

We now break the calculation into two pieces, as explained in Section 4.2. However, Peter and Meyer-ter-Vehn use the term “individual” rather than “collisional” since there are no collisions in a purely dielectric formulation. In this

calculation,

$$L_{div} = -\frac{4}{\pi} \int_1^{K_{max}} K^3 dK \int_{-V_t/\sqrt{2}}^{V_t/\sqrt{2}} ds \frac{sY(s)}{(2K^2 - X(s))^2 + Y^2(s)} \quad (4.40)$$

and

$$L_{dp.w.} = -\frac{4}{\pi} \int_0^1 K^3 dK \int_{-V_t/\sqrt{2}}^{V_t/\sqrt{2}} ds \frac{sY(s)}{(2K^2 - X(s))^2 + Y^2(s)}, \quad (4.41)$$

where splitting the integral for $K > 1$ and $K < 1$ is arbitrary.

To perform the integration for the plasma wave piece, we need to look at the expansions for X and Y ,

$$X(s) = \frac{1}{s^2} + \frac{3}{2} \frac{1}{s^4} + \frac{15}{4} \frac{1}{s^6} + \dots \quad (4.42)$$

and

$$Y(s) = -2\sqrt{\pi} s e^{-s^2}. \quad (4.43)$$

It is now somewhat apparent that, for $V_t \gg 1$, the integration along the real μ axis comes close to a pole in the integrand. We use the fact that

$$\lim_{Y \rightarrow 0} \frac{Yg(s)}{f^2(s) + Y^2} = \pi \delta(s - s_0) \frac{\text{sign}(Y)g(s_0)}{|f'(s_0)|}, \quad (4.44)$$

where s_0 is a zero of f . For sufficiently small K (we used $K < 1$ when we split the overall integral into equations 4.40 and 4.41), $f(s) = 2K^2 - X(s) \approx 2K^2 - \frac{1}{s^2}$ and $s_0 = \pm \frac{1}{\sqrt{2K}}$. Since $|s| < V_t/\sqrt{2}$, the path of integration only passes by the pole when $K > \frac{1}{V_t}$. The physical meaning of this restriction is that there are no interactions with particles farther away than the dynamic screening length, L_t . In the end,

$$L_{dp.w.} \approx -\frac{4}{\pi} \int_{1/V_t}^1 K^3 dK 2 \int_0^{V_t/\sqrt{2}} ds \pi \delta(s - s_0) \frac{-|s_0|}{|2/s_0^3|}. \quad (4.45)$$

Therefore,

$$\begin{aligned} L_{dp.w.} &\approx \frac{4}{\pi} \int_{1/V_t}^1 K^3 dK \frac{\pi}{4K^4} \\ &\approx \int_{1/V_t}^1 \frac{dK}{K} \\ &\approx \ln(V_t). \end{aligned} \quad (4.46)$$

For the individual piece, we make the (rather bad) approximation that, for $K > 1$ and $V_t \gg 1$,

$$\frac{sY(s)}{(2K^2 - X(s))^2 + Y^2(s)} \approx \frac{sY(s)}{4K^4}. \quad (4.47)$$

Then,

$$L_{d\text{indiv}} \approx -\frac{4}{\pi} \int_1^{K_{max}} K^3 dK 2 \int_0^{V_t/\sqrt{2}} ds \frac{sY(s)}{4K^4}. \quad (4.48)$$

Since $V_t \gg 1$, we can safely extend the upper limit of integration on s , so that

$$L_{d\text{indiv}} \approx -\frac{4}{\pi} \int_1^{K_{max}} K^3 dK 2 \int_0^\infty ds \frac{s \left(-2\sqrt{\pi} s e^{-s^2} \right)}{4K^4}. \quad (4.49)$$

Performing the s integral,

$$\begin{aligned} L_{d\text{indiv}} &\approx -\frac{4}{\pi} \frac{q_t^2}{\lambda_{De}^2} \int_1^{K_{max}} \frac{-\pi dK}{4K} \\ &\approx \int_1^{K_{max}} \frac{dK}{K} \\ &\approx \ln(K_{max}). \end{aligned} \quad (4.50)$$

We add the two stopping pieces, Equations 4.46 and 4.50, together to get

$$\begin{aligned} L_d &\approx \int_{1/V_t}^{K_{max}} \frac{dK}{K} \\ &\approx \ln(K_{max} V_t) \\ &\approx \ln \left(\frac{2\pi N_D V_t^3}{|Z_t|} \right). \end{aligned} \quad (4.51)$$

Or, in non-normalized variables,

$$L_d \approx \ln \left(\frac{m_e v_t^3}{2\omega_{pe} q_t e} \right), \quad (4.52)$$

which is Equation 4.5 when $m_t \gg m_e$. This formulation clearly shows that the integration of $1/K$ is the source of the logarithmic dependence, and that the $1/V_t$ cutoff is not arbitrarily imposed but derived, and occurs due to plasma screening. It is also worth remembering (as is discussed further in the next section) that, for large K_{max} , Equation 4.31 makes no sense. It is derived under the straight-line or differential approximation, which is clearly not reasonable for small impact parameters, under which large angle scattering and large exchanges of energy occur.

4.4.2 Typical (Collisional) Derivation

We now follow the typical calculation for the collisional piece. Using the formula for binary collision energy transfer [GPS01],

$$\frac{\Delta\mathcal{E}_c}{\mathcal{E}} = \frac{1}{1 + \left(\frac{2b}{d_{ca}}\right)^2} \frac{4m_t m_e}{(m_t + m_e)^2}, \quad (4.53)$$

where d_{ca} is given in Equation 4.4. The use of Equation 4.53 is the Landau-Boltzmann point of view, while the previous section is the Lenard-Bolescu point of view.

Now, assume that the test charge is moving through the plasma in some direction (we choose \hat{x}). As the test charge moves a distance dx , we imagine that it goes through collisions with particles in the transverse direction out to a maximum distance b_{max} . In differential form, we write this energy transfer as

$$-d\mathcal{E}_c = dx dy dz n_0 \Delta\mathcal{E}_c. \quad (4.54)$$

We note that this differential form makes no sense when $\Delta\mathcal{E}_c \approx \mathcal{E}$, as occurs for $b \rightarrow 0$ and $m_t = m_e$. Integrating over the transverse direction,

$$-\left(\frac{d\mathcal{E}}{dx}\right)_c = \int \int dy dz n_0 \Delta E_c. \quad (4.55)$$

We make the change to cylindrical coordinates, so

$$-\left(\frac{d\mathcal{E}}{dx}\right)_c = \int_0^{2\pi} d\phi \int_{b_{min}}^{b_{max}} b db \Delta E_c, \quad (4.56)$$

or

$$\begin{aligned}
-\left(\frac{1}{\mathcal{E}} \frac{d\mathcal{E}}{dx}\right)_c &= 2\pi n_0 \int_{b_{min}}^{b_{max}} \frac{bdb}{1 + \left(\frac{2b}{d_{ca}}\right)^2} \frac{4m_t m_e}{(m_t + m_e)^2} \\
&= 2\pi n_0 d_{ca}^2 \frac{4m_t m_e}{(m_t + m_e)^2} \int_{b_{min}}^{b_{max}} \frac{bdb}{d_{ca}^2 + 4b^2} \\
&= 2\pi n_0 \frac{4q_t^2 e^2 (m_t + m_e)^2}{v_t^4 (m_t m_e)^2} \frac{4m_t m_e}{(m_t + m_e)^2} \int_{b_{min}}^{b_{max}} \frac{bdb}{d_{ca}^2 + 4b^2} \\
&= 8\pi n_0 \frac{4q_t^2 e^2}{v_t^4 m_t m_e} \int_{b_{min}}^{b_{max}} \frac{bdb}{d_{ca}^2 + 4b^2} \\
&= 4 \frac{q_t^2 \omega_{pe}^2}{v_t^2} \frac{2}{m_t v_t^2} \int_{b_{min}}^{b_{max}} \frac{bdb}{d_{ca}^2 + 4b^2} \tag{4.57}
\end{aligned}$$

This integral converges as $b_{min} \rightarrow 0$, but logarithmically diverges as $b_{max} \rightarrow \infty$. Performing the integral and writing the stopping power in terms of L_d ,

$$L_{dc} = \frac{1}{2} \ln \left(\frac{d_{ca}^2 + 4b_{max}^2}{d_{ca}^2 + 4b_{min}^2} \right). \tag{4.58}$$

If $b_{max} \gg d_{ca}$ (a reasonable assumption) and $b_{min} \rightarrow 0$,

$$\begin{aligned}
L_{dc} &\approx \frac{1}{2} \ln \left(\frac{4b_{max}^2}{d_{ca}^2} \right) \\
&\approx \ln \left(\frac{2b_{max}}{d_{ca}} \right). \tag{4.59}
\end{aligned}$$

Interestingly, the result is well-behaved as $b_{min} \rightarrow 0$ even though the concept of a $d\mathcal{E}/ds$ makes no sense.

For the plasma wave piece, we use Equation 4.39,

$$L_d = -\frac{4}{\pi} \int_0^{K_{max}} K^3 dK \int_{-V_t/\sqrt{2}}^{V_t/\sqrt{2}} ds \frac{sY(s)}{(2K^2 - X(s))^2 + Y^2(s)}. \tag{4.60}$$

We follow a process similar to that for the plasma wave piece in Section 4.4.1, except we do not split the integral at $K = 1$ and we set $K_{max} = \lambda_{De}/b_{max}$. Using the approximation in Equation 4.44,

$$\begin{aligned}
L_{dp.w.} &\approx \frac{4}{\pi} \int_{1/V_t}^{K_{max}} K^3 dK \frac{\pi}{4K^4} \\
&\approx \ln(K_{max} V_t). \tag{4.61}
\end{aligned}$$

In non-normalized variables,

$$L_{dp.w.} \approx \ln \left(\frac{v_t}{b_{max}\omega_{pe}} \right). \quad (4.62)$$

The final stopping power once again becomes

$$\begin{aligned} L_d &= L_{dp.w.} + L_{dc} \\ &\approx \ln \left(\frac{v_t}{b_{max}\omega_{pe}} \right) + \ln \left(\frac{2b_{max}}{d_{ca}} \right) \\ &\approx \ln \left(\frac{2v_t}{d_{ca}\omega_{pe}} \right) \\ &\approx \ln \left(\frac{v_t^3 m_t m_e}{q_t e (m_t + m_e) \omega_{pe}} \right). \end{aligned} \quad (4.63)$$

For ion stopping ($m_t \gg m_e$), this formula reduces to the Equation 4.5, but differs by a factor of two inside the logarithm. We find in our comparisons in Section 4.7.1 that this version of the Bohr formula agrees best with molecular dynamics simulations.

4.5 Wake and Collision Contributions

Now that we have covered the detailed Bohr derivation, we can attempt to get an idea about the collisional and wake contributions to the stopping power. Rewriting Equation 4.46 in terms of impact parameter, converting to non-normalized units, and using $k = 1/b$,

$$L_{dp.w.} \approx \int_{b_{max}}^{v_t/\omega_{pe}} \frac{db}{b}. \quad (4.64)$$

Normalizing so that $B = b/d_{ca}$, the the wake portion is,

$$L_{dp.w.} \approx \int_{B_{max}}^{\frac{v_t}{\omega_{pe} d_{ca}}} \frac{dB}{B} \quad (4.65)$$

and the collisional portion—from Equation 4.57—is,

$$L_{dc} \approx \int_{B_{min}}^{B_{max}} dB \frac{4B}{1 + 4B^2}. \quad (4.66)$$




Figure 4.1: The integrands of the so-called wake (blue) and collisional (red) portions of Bohr stopping power as a function of impact parameter. Starting around $b \sim L_t$, the full dielectric integrand includes collective screening, and will fall below the collisional result, thereby removing the logarithmic large- b divergence.

We plot the integrands from the wake and collisional portions as a function of impact parameter in Figure 4.1. The two diverge for $b \lesssim 2d_{ca}$ but overlap for $b \gtrsim 2d_{ca}$. The plot demonstrates that there is no clear transition point from one formulation to the other. There is a regime where the two overlap and both apply.

4.6 The QED Stopping Formula

During the past decade, agreement has been reached concerning the stopping and scattering of a single relativistic electron in a plasma. The calculation is based upon QED and uses the Møller cross-section [LP06, SB08, ASD09]. A modification to the cross-section for materials, which also applies to plasmas,

was originally discussed in ICRU Report 37 [RM84], and was used in References [SB08, ASD09]. The final version of the calculation, including scattering formulas and without mistakes, was published by Atzeni et. al. [ASD09]. We cover the detailed stopping power derivation in this Section.

The collision term is calculated using the Møller cross-section,

$$\frac{d\sigma}{dW} = \frac{2\pi e^4}{m_e c^2 \beta^2} \left[\frac{1}{W^2} + \frac{1}{(\mathcal{E} - W)^2} + \left(\frac{\gamma - 1}{\gamma \mathcal{E}} \right)^2 - \frac{2\gamma - 1}{\gamma^2 W (\mathcal{E} - W)} \right], \quad (4.67)$$

where W is the energy transferred during a collision and \mathcal{E} is the kinetic energy of the incident electron. The first two terms in the brackets correspond to projectile and target Rutherford scattering, respectively. The third term comes from electron spin. The last term takes into account exchange effects since the electrons are indistinguishable.

The collisional stopping power is then given by

$$\left(\frac{d\mathcal{E}}{ds} \right)_c = -n_e \int_{W_c}^{\mathcal{E}/2} W (d\sigma/dW) dW, \quad (4.68)$$

where W_c is the minimum energy transferred during a collision. The upper limit of integration is $\mathcal{E}/2$ due to symmetry, so the electron with the greater energy after the collision is considered the projectile. Performing the integration,

$$\begin{aligned} \left(\frac{d\mathcal{E}}{ds} \right)_c = & -n_e \frac{2\pi e^4}{m_e c^2 \beta^2} \left\{ \left[\ln \left(\frac{\mathcal{E}/2}{W_c} \right) \right] \right. \\ & + \left[2 + \ln(\mathcal{E}/2) - \frac{\mathcal{E}}{\mathcal{E} - W_c} - \ln(\mathcal{E} - W_c) \right] \\ & + \left[\left(\frac{\gamma - 1}{\gamma \mathcal{E}} \right)^2 \frac{(\mathcal{E}/2)^2 - W_c^2}{2} \right] \\ & \left. - \left[\frac{2\gamma - 1}{\gamma^2} (-\ln(\mathcal{E}/2) + \ln(\mathcal{E} - W_c)) \right] \right\}. \quad (4.69) \end{aligned}$$

Now assume that $\mathcal{E} \gg W_c$ and write the stopping power in terms of L_d , so that

$$\begin{aligned}
L_{dc} = \frac{1}{2} & \left\{ \left[\ln \left(\frac{\mathcal{E}}{W_c} \right) - \ln(2) \right] \right. \\
& + [2 + \ln(\mathcal{E}) - \ln(2) - 1 - \ln(\mathcal{E})] \\
& + \left[\left(\frac{\gamma - 1}{\gamma \mathcal{E}} \right)^2 \frac{\mathcal{E}^2}{8} \right] \\
& \left. + \left[\frac{2\gamma - 1}{\gamma^2} (\ln(\mathcal{E}) - \ln(2) - \ln(\mathcal{E})) \right] \right\}. \tag{4.70}
\end{aligned}$$

All that's left is to bring the 1/2 inside the braces, perform the obvious cancellations, and group like terms. Expanding everything out before the final grouping gives,

$$\begin{aligned}
L_{dc} = & \left\{ \left[\frac{1}{2} \ln \left(\frac{\mathcal{E}}{W_c} \right) - \frac{\ln(2)}{2} \right] + \left[\frac{1}{2} - \frac{\ln(2)}{2} \right] \right. \\
& \left. + \left[\frac{1}{16} - \frac{1}{8\gamma} + \frac{1}{16\gamma^2} \right] - \left[\ln(2) \left(\frac{1}{\gamma} - \frac{1}{2\gamma^2} \right) \right] \right\}. \tag{4.71}
\end{aligned}$$

The final expression for the collisional component is then

$$L_{dc} = \left[\frac{1}{2} \ln \left(\frac{\mathcal{E}}{W_c} \right) + \frac{9}{16} - \ln 2 + \frac{\ln 2 + 1/8}{2\gamma^2} - \frac{\ln 2 + 1/8}{\gamma} \right]. \tag{4.72}$$

The dielectric component in this case has transverse (Bethe's formula) and longitudinal (density effect) pieces. The derivation is long and can be found in many places. See, e.g. chapter 13 of Reference [Jac99]. Note, however, that the derivation is in the cold plasma limit, which is identical to that in Section 4.3.

The result is

$$L_{d.p.w.} = \frac{1}{2} \ln \left(\frac{2m_e c^2 \beta^2 W_c}{\hbar^2 \omega_{pe}^2} \right). \tag{4.73}$$

Adding the dielectric component to the collisional component gives

$$L_d = \left[\frac{1}{2} \ln \left(\frac{2m_e c^2 \beta^2 \mathcal{E}}{\hbar^2 \omega_{pe}^2} \right) + \frac{9}{16} - \ln 2 + \frac{\ln 2 + 1/8}{2\gamma^2} - \frac{\ln 2 + 1/8}{\gamma} \right], \tag{4.74}$$

which is often written as

$$L_d = \left[\ln \left(\frac{\sqrt{m_e c^2 \beta^2 \mathcal{E}}}{\hbar \omega_{pe}} \right) + \frac{9}{16} - \frac{\ln 2}{2} + \frac{\ln 2 + 1/8}{2\gamma^2} - \frac{\ln 2 + 1/8}{\gamma} \right]. \tag{4.75}$$

Formula 4.75 is currently considered the most accurate electron stopping power formula for Fast Ignition, and is used in many simulations today. However, as we demonstrate in Chapter 6, correlations between beam electrons can cause the overall beam stopping power to be greater than the stopping power given by this formula.

4.7 Comparison with Simulations

We compare the single-particle stopping powers with four different simulation codes, the molecular dynamics (MD) code ddcMD, the electrostatic PIC code BEPS, the fully electromagnetic finite-difference PIC code OSIRIS, and the quasi-static PIC code QuickPIC. We simulate non-relativistic electron and ion stopping powers using ddcMD and BEPS, and relativistic electron stopping powers using OSIRIS and QuickPIC. The codes ddcMD and BEPS are used in Chapter 3, and are described in detail there. We describe OSIRIS in this Section, and a description of QuickPIC is in Appendix A.

We compare the results from all four codes with theories. The Bohr formula from Equation 4.63 agrees well with results from ddcMD when the projectile speed $v_t \gg v_{th}$. Computational requirements limited the cell width in BEPS, with the smallest $\Delta = 0.125\lambda_{De}$. The BEPS results do not agree with the Bohr formula, but do agree with stopping powers we calculate using the wake script described in Section 3.5. QuickPIC stopping powers for larger cell widths (Δ s) agree with the stopping power calculated using fluid theory to within a constant factor (they should agree if the effect of finite size particles is calculated exactly), but unlike fluid theory, they saturate once the cell width shrinks beyond a certain point. OSIRIS has the most difficulty due to numerical Čerenkov, which causes the projectiles to slow down significantly when passing through free space for small cell widths. We note that new field solvers may eliminate this issue [LLT13, MTE14],

Electrostatic Simulation Parameters	
n_0	$1.03 \times 10^{20} \text{ cm}^{-3}$
T_e	10.8846 eV
λ_{De}	24.2 Å
$n_0 \lambda_{De}^3$	1.45
Coupling Parameter, Γ	0.06
m_p/m_e	1824
ddcMD Box Dimensions	$267\lambda_{De} \times 267\lambda_{De} \times 267\lambda_{De}$
ddcMD dt	$2.85 \times 10^{-5} \omega_{pe}^{-1}$
BEPS Box Dimensions	$256\lambda_{De} \times 256\lambda_{De} \times 256\lambda_{De}$
BEPS dt	$(0.2, 0.05, 0.01, 0.01) \omega_{pe}^{-1}$ for $\Delta = (1, 1/2, 1/4, 1/8) \lambda_{De}$

Table 4.1: The parameters for the electrostatic single-particle stopping power simulations using ddcMD and BEPS.

pointing to directions for future work. These results are presented in detail in the following subsections.

4.7.1 Comparisons with ddcMD and BEPS

As described in Chapter 3, ddcMD and BEPS are both periodic, electrostatic, particle codes. The primary difference is that BEPS is a PIC code, in which particles only interact via fields calculated on a mesh, while ddcMD uses the particle-particle particle-mesh (PPPM) method to resolve nearest-neighbor interactions. We compare results from both with the Bohr formula (Equations 4.5 and 4.63), and compare the BEPS results with stopping power calculated using the wake script used in Section 3.5. Both BEPS and ddcMD use electrostatic field solvers. We will study the stopping of relativistic electrons in the next two chapters.

The BEPS and ddcMD stopping power simulations are performed using an

electron plasma with parameters given in Table 4.1. Simulations from both codes use a smooth neutralizing positive background. In the BEPS simulations, a single PIC particle represents a single physical particle, which avoids the enhanced stopping that sometimes occurs in PIC simulations [MTE14]. Note that the proton-electron mass ratio is a legacy value from early ddcMD stopping power simulations, and differs slightly from the correct physical value of 1836.15. This difference results in a very small change in the stopping power. The ddcMD simulations also use a Langevin thermostat [DH09] with a decay time of $350.5\omega_{pe}^{-1}$ to prevent the plasma from heating. We also use small values of $n_0\lambda_{De}^3$. This makes the agreement between BEPS and ddcMD greater (as expected). For larger but still small values of $n_0\lambda_{De}^3$, there should be better agreement, showing that BEPS can be used to study collisions.

The simulations use 27 projectiles spaced maximally on a cubic lattice. The projectiles are always negatively charged to avoid numerical problems with an attractive Coulomb singularity. All the projectiles move in the direction $\hat{v} = \langle 1, \sqrt{\varphi}, \varphi \rangle$, where $\varphi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. The direction is chosen to minimize the chance of the projectiles traveling through each others' wakes as they loop around the periodic box.

Stopping powers from the simulations are calculated using linear regression of the quantity $\frac{1}{2}E_tv_t^2$ with s_t for each projectile, where E_t is the measured energy of the projectile, v_t its measured speed, and s_t the distance it has traveled since the beginning of the simulation. This linear regression gives us (after some substitutions involving E_t and v_t) $v_t^2 \frac{dE_t}{ds_t}$, which is the stopping power in the units we plot in this section. We then use the stopping power from all 27 projectiles in each simulation to calculate the average stopping power and the standard deviation.

The results of ddcMD simulations with anti-(zinc, neon, and hydrogen) nuclei as projectiles, which have $Z=(-30, -10, \text{ and } -1)$, respectively, are in Figures 4.2 and 4.4. The projectiles start with a speed of $V_0 = 11.28v_{th}$. The projectile energies

agree well with a calculation using the the Bohr stopping power in Equation 4.63, which is plotted as a dashed-dotted black line in Figure 4.2. In particular, in the anti-zinc case, Equation 4.63 produces a result inside the simulation error bars, while Equation 4.5 does not. At various points in Figure 4.2 there are sudden drops in the energy of the simulation projectiles, which are caused by close encounters with plasma electrons and contribute to the stopping. As we will see next, these drops do not occur in PIC simulations. For these simulations (with small $n_0\lambda_{De}^3$) large angle scattering events (in the rest mass frame) occur.

The results of BEPS simulations with anti-zinc nuclei and anti-protons as projectiles are shown in Figures 4.3 and 4.4. Stopping powers for simulations with anti-neon nuclei projectiles are also in Figure 4.4. The projectiles start with a speed of $V_0 = 11.25v_{th}$, which does not quite match with the ddcMD simulations. These BEPS simulations face the issue of wake build-up in the box because they lack significant collisional damping (of the wake) and the code does not have a Langevin thermostat. The wake build-up eventually affects the energy dissipation of the projectiles, as seen by the wiggles in the curves for the anti-zinc nuclei in Figure 4.3. We therefore neglect data after $s_t = 80\lambda_{De}v_0/v_{th}$ when calculating the stopping power in the BEPS simulations. Reducing the cell width increases the collisionality, which increases the damping and reduces the wiggles.

During the time before the wake build-up affects the projectile energy in the BEPS simulations, the projectile energies agree well with a calculation using a constant stopping power from the wake script, which we plot as dashed-dotted black lines in Figure 4.3. We use a box with dimensions $64\lambda_{De} \times 64\lambda_{De} \times 128\lambda_{De}$ for the wake script stopping power calculation. The projectile starts at $\vec{x}_0 = (32\lambda_{De}, 32\lambda_{De}, 16\lambda_{De})$. We assign the projectile a velocity $\vec{v} = 11.25v_{th}\hat{z}$ and perform the calculation at $t = 5\omega_{pe}^{-1}$, giving the projectile a final position $\vec{x} = (32\lambda_{De}, 32\lambda_{De}, 72.23\lambda_{De})$. We calculate the force on the projectile by interpolating the E_z field to the particle's position using the algorithm from BEPS.




Figure 4.2: Energy vs. distance traveled in ddcMD simulations using as projectiles anti-zinc nuclei (top), anti-neon nuclei (middle), and anti-protons (bottom). The dashed-dotted black line is the energy calculated using the Bohr formula (Equation 4.63), while the solid lines are the 27 simulation projectiles.




Figure 4.3: Energy vs. distance traveled in BEPS simulations using as projectiles anti-zinc nuclei in the top two plots and anti-protons in the bottom plot. The cell-width is λ_{De} in the top plot and $\lambda_{De}/8$ in the bottom two plots. The dashed-dotted black line is the stopping power calculated using the wake script while the solid lines are the 27 simulation projectiles.

As the cell width decreases in the the BEPS simulations, the stopping power increases and approaches the Bohr stopping power, as shown in Figure 4.4. However, when we compare Figure 4.3 with Figure 4.2, we notice that the projectiles in the BEPS simulations do not undergo the sudden drops in energy that the projectiles in ddcMD do, suggesting that the PIC finite-size particles are still smoothing collisional contributions to the stopping power for small impact parameters even at the smallest cell widths. Stated another way, collisions with impact parameters less than a cell width still occur. This hypothesis is also supported by the excellent agreement between simulation and the wake script with cell widths of $0.125\lambda_{De}$. Due to computing resource constraints, we are unable to study the effects of cell widths below $0.125\lambda_{De}$, where we may begin seeing the stopping power in BEPS diverge from that in the wake script and converge closer to ddcMD.

We examine the dependence of stopping power on speed in Figure 4.5 for anti-zinc nuclei in both ddcMD and BEPS. The agreement between the Bohr formula and ddcMD holds up well until the speed is below $6v_{th}$. The divergence is expected, as the logarithm diverges as its argument approaches zero, which is unphysical. The agreement between BEPS and the wake script, however, remains impressive across the three test cases. Vlasov theory clearly describes what is occurring in BEPS quite well down to relatively low speeds.

Because of their low mass, electrons behave differently to ions in the stopping power simulations. As we show in Figure 4.6 for electrons with $V_0 = 11.28v_{th}$ and $11.25v_{th}$ in ddcMD and BEPS, respectively, the energies spread out much more rapidly, which is a heating effect. The spread is much greater in ddcMD, likely due to collisions, which are a means of energy transfer between the projectiles and the plasma. This heating effect makes comparing with stopping power formulas more difficult, since there is a significant spread in speed. When we calculate the energy for plots using the Bohr formula, we use the average speed at each step. For the wake script, we only use the initial stopping power due to the time




Figure 4.4: The stopping powers measured in ddcMD (blue circles) and BEPS (red circles) simulations for anti-protons (left), anti-neon nuclei (middle), and anti-zinc nuclei (right). The ddcMD results are plotted with a cell width of zero for convenience. Projectiles in ddcMD start with $V_0 = 11.28v_{th}$ while those in BEPS start with $V_0 = 11.25v_{th}$. Stopping powers calculated using the Bohr formula and the wake script (red squares) are plotted for comparison with the ddcMD and BEPS results, respectively. Bohr 1 (blue squares) uses Equation 4.63 while Bohr 2 (blue diamonds) uses Equation 4.5, with the term inside the logarithm in Bohr 1 being twice that of Bohr 2.




Figure 4.5: The stopping powers measured in ddcMD and BEPS simulations for anti-zinc nuclei starting at various speeds along with comparisons to the Bohr stopping power from Equation 4.63 and the wake script, respectively. The cell width in BEPS is $0.125\lambda_{De}$. The ddcMD stopping powers are blue circles while the Bohr stopping power is a solid blue line. The BEPS stopping powers are red circles while the stopping power from the wake script is a solid red line.

required to perform the calculation. A more direct comparison of the stopping powers, along with results for more BEPS cell widths, is in Figure 4.8.

The other factor we must consider in electron stopping is angular scattering. We did not concern ourselves with angular scattering of ions since they are much more massive than electrons, and their deflection in electron plasmas is negligible. In Figure 4.7 we plot the deflection of electrons from their initial trajectories in the ddcMD and BEPS simulations mentioned previously. The projectiles clearly undergo more deflection in ddcMD, because it resolves collisions better than BEPS. A more direct comparison of the angular deflections, along with results for more BEPS cell widths, is also in Figure 4.8.

The comparisons in Figure 4.8 demonstrate that the angular scattering in BEPS approaches that in ddcMD as we decrease the cell width, which is due to the fact that we resolve collisions better as we decrease the BEPS cell width. The stopping power does not demonstrate a similar trend, because heating of the projectiles obscures the stopping. If we were able to further reduce the cell width, a trend might manifest itself.

The agreement with theoretical stopping calculations is best when collisions do not play a large role, minimizing heating effects. At the smallest BEPS cell width and in ddcMD, the agreement is acceptable, but not very good. The story may change with speed, and relativistic effects may help minimize heating, but we did not have the opportunity to study heating in detail. These comparisons simply demonstrate the complexity of the stopping power problem, and that thermal effects must be considered before attempting to draw solid conclusions about electron beam stopping.

The simulations and comparisons demonstrate that we understand classical ion stopping power in an electron plasma in the electrostatic regime for fast projectiles ($v_t \gg v_{th}$) very well, and that understanding is extended to slow projectiles in PIC. In these situations, we can explain the basic Coulomb stopping power using the




Figure 4.6: Energy vs. distance traveled for electrons in ddcMD and BEPS simulations are in the top and bottom plots, respectively. The electrons start with $V_0 = 11.28v_{th}$ and $11.25v_{th}$ in ddcMD and BEPS, respectively. The BEPS simulation uses a cell width of $0.125\lambda_{De}$. The dashed-dotted black line is the energy calculated using the Bohr formula and the wake script for comparison with ddcMD and BEPS, respectively.




Figure 4.7: Perpendicular vs. parallel distance traveled for electrons in ddcMD and BEPS are in the top and bottom plots, respectively, for electrons starting with $V_0 = 11.28v_{th}$ and $11.25v_{th}$, respectively. The BEPS simulation uses a cell width of $0.125\lambda_{De}$. s_{\parallel} is the distance traveled along the electron's initial trajectory, while s_{\perp} is the distance traveled perpendicular to it.




Figure 4.8: Stopping power and angular displacement of electrons in ddcMD (blue circles) and BEPS (red circles) simulations. The ddcMD results are plotted with a cell width of zero for convenience. Projectiles in ddcMD start with $V_0 = 11.28v_{th}$ while those in BEPS start with $V_0 = 11.25v_{th}$. Stopping powers calculated using the Bohr formula and the wake script (red squares) are plotted for comparison with the ddcMD and BEPS results, respectively. Bohr 1 (blue squares) uses Equation 4.63 while Bohr 2 (blue diamonds) uses Equation 4.5, with the term inside the logarithm in Bohr 1 being twice that of Bohr 2.

Bohr stopping power formula and PIC stopping power using our wake script. As the mass decreases and thermal effects become important, the situation becomes more complicated, because the projectiles can thermalize with the background plasma and angular scattering becomes important. We have not studied this thermalization effect in detail, and are uncertain about its importance in the relativistic regime. These thermalization effects should be studied in more detail before drawing solid conclusions about electron beam stopping.

4.7.2 Comparison with OSIRIS

As mentioned earlier, OSIRIS is a finite-difference time-domain (FDTD) electromagnetic PIC code, and therefore suffers from numerical Čerenkov radiation. This radiation is produced when particles approach the speed of light, as plotted in Figure 4.9. It is caused by particles exceeding the speed of electromagnetic wave on the mesh (the mesh acts like a slow wave structure). The finite-difference operators modify the electromagnetic dispersion relation and cause light to propagate slower than c for some frequencies in some directions. As the particle radiates, it slows down due to conservation of energy. The problem becomes worse as the cell width decreases, rendering OSIRIS largely useless as a tool for our stopping power studies unless improved field solvers can be developed. Spectral electromagnetic PIC codes, such as the electromagnetic version of BEPS, do not suffer from this problem due to their high order of accuracy, and QuickPIC does not suffer from it due to its quasi-static approximations.

Methods for mitigating numerical Čerenkov radiation include smoothing the fields and using higher-order solvers. As we saw in Section 2.3.3, smoothing is used to damp certain \vec{k} modes, so we can take advantage of that fact. However, our attempts to improve the smoother beyond the five-pass one described in Section 2.3.3 did not reduce the stopping power any further.

OSIRIS Simulation Parameters	
n_0	10^{26} cm^{-3}
T_e	0 (cold)
Box Dimensions	$20.7575 \frac{c}{\omega_{pe}} \times 15.568125 \frac{c}{\omega_{pe}} \times 15.568125 \frac{c}{\omega_{pe}}$
Box Cells	$512 \times 384 \times 384$ for $\Delta_0 = 0.0405 \frac{c}{\omega_{pe}}$
dt	$(0.01, 0.01, 0.005) \omega_{pe}^{-1}$ for $\Delta = (1, 1/2, 1/4) \Delta_0$

Table 4.2: The parameters for the OSIRIS stopping power simulations.

A much greater benefit comes from using a more accurate solver. We tested the extended stencil described by Greenwood et. al. in Reference [GCL04] with coefficients that produce a curl operator accurate to fourth order. For comparison, the standard Yee solver is second-order accurate. The new solver reduces the free-space stopping power by about a factor of seven, but that is not enough, as the numerical Čerenkov radiation still dominates the stopping power in simulations using smaller cell widths, as shown in Figure 4.10.

The parameters for the simulations we performed with OSIRIS are in Table 4.2. As in the BEPS simulations, a single PIC particle represents a single physical particle. We perform simulations with cell widths $\Delta = \Delta_0, \Delta_0/2,$ and $\Delta_0/4,$ where $\Delta_0 = 0.0405 \frac{c}{\omega_{pe}}$ is the plasma electron spacing at $n_0 = 10^{26} \text{ cm}^{-3}$. The box has periodic boundary conditions in the transverse directions and a Galilean moving window at c in the \hat{z} -direction. We smooth both the fields and the current using the five-pass smoother mentioned above. The projectile is a single electron, which starts at rest and is uniformly accelerated to $p_z = 10m_e c$ across one hundred time-steps. Since the electric field is zero at $t = 0$, the projectile leaves behind a fixed neutralizing charge as it moves away from its starting point.

Figure 4.10 shows the results of single-particle stopping power simulations using $\Delta = \Delta_0/4$ with both the standard Yee solver and the fourth-order stencil solver. To demonstrate the effect of the numerical Čerenkov radiation, we plot




Figure 4.9: The E_z field of an electron's wake at $t = 20\omega_{pe}^{-1}$ in an OSIRIS simulation with $\Delta = \Delta_0/2$ using the Yee solver. The slice is taken down the middle of the box in y . Numerical Čerenkov radiation, caused by the electron exceeding the propagation speed of light on the mesh, is radiated off to the sides.

the stopping power from simulations with and without a background plasma. The plots show stopping power vs. momentum, so the particle starts on the right of the momentum axis. The sudden drop in stopping power at the beginning of the simulation is a residual effect due to the acceleration of the projectile from rest and its movement away from its neutralizing charge. After the stopping power levels off, the numerical Čerenkov radiation clearly dominates for both solvers.

The fourth order stencil solver has an additional complication of inconsistent stopping power. Notice the curvature in the plot, which is not present with the Yee solver. The stopping power oscillates. At this time, we do not know the reason for this oscillation. While it is problematic for our stopping power simulations, it is likely unnoticeable in many applications which can derive a benefit from reduced numerical Čerenkov radiation.

Due to the problem with numerical Čerenkov radiation, OSIRIS currently is not suitable for our studies. Instead, we have elected to use QuickPIC, whose quasi-static approximation does not allow for numerical Čerenkov radiation. Recent work has demonstrated that a new solver mitigates the problem [LLT13], which opens up opportunities to study correlated stopping in more detail, by taking into account beam divergence and angular scattering. For now, those effects are left as areas for future research.

4.7.3 Comparison with QuickPIC

QuickPIC uses a quasi-static approximation, which assumes that the beam evolves on a time-scale much slower than the time required for the beam to pass by a plasma particle. The primary requirement is that the beam must be traveling close to the speed of light (in practice, $\gamma \gtrsim 5$). Under these assumptions, we can take advantage of the slow wake evolution to take much larger larger time-steps than in typical PIC simulations. For more details about QuickPIC, see Appendix




Figure 4.10: The stopping powers measured in OSIRIS simulations using $\Delta = \Delta_0/4$ with the traditional Yee solver and the fourth-order accurate extended stencil. The red line is the stopping power in a simulation with a plasma, while the blue line is from a simulation without a plasma.

QuickPIC Simulation Parameters	
n_0	10^{26} cm^{-3}
T_e	0.0 eV
Box Dimensions	$10.3787 \frac{c}{\omega_{pe}} \times 10.3787 \frac{c}{\omega_{pe}} \times 0.3243 \frac{c}{\omega_{pe}}$
Box Cells	$256 \times 256 \times 8$ for $\Delta = 0.0405 \frac{c}{\omega_{pe}}$
dt	ω_{pe}^{-1} for all cell widths

Table 4.3: The parameters for the QuickPIC stopping power simulations.

A.

The parameters for the QuickPIC simulations are in Table 4.3. The box has a moving window in the \hat{z} -direction, as in OSIRIS, but uses conducting boundary conditions with specular reflection for particles in the transverse direction. The electron projectile is initialized with $p_z = 10m_e c$. Unlike in the OSIRIS simulations, we do not need to worry about residual effects from the initialization due to the nature of the quasi-static implementation in QuickPIC. We start with a cell-width $\Delta_0 = 0.0405 \frac{c}{\omega_{pe}}$ and shrink it by a factor of two in each successive simulation until the final simulation uses $\Delta = \frac{1}{64} \Delta_0$.

As in OSIRIS and BEPS, we have chosen to have a PIC particle represent a physical particle. However, this representation only applies in the transverse direction. QuickPIC uses a 2D plasma slice that traverses the box during each 3D time-step and passes through every cell in z , as described in Appendix A. Therefore, the concept of the number of cells per particle in z does not truly apply to QuickPIC.

Figure 4.11 demonstrates that the single-particle stopping power in QuickPIC agrees well with the fluid wake stopping power from Equation 4.30 until $\Delta < \Delta_0/4$, differing only by a factor of two in the logarithm. The agreement may improve by taking into account the particle shape and smoothing in more detail. We also plot the Bethe, QED, and Bohr stopping formulas for comparison, which




Figure 4.11: The stopping power in the QuickPIC simulations (blue with circles) compared with the fluid wake stopping power from Equation 4.30 (black with squares), the Bethe formula in Equation 4.6 (dashed green), the QED formula in Equation 4.75 (dashed red), and the Bohr formula in Equation 4.63 (dashed magenta).

are all far above the simulated QuickPIC stopping power. The Bohr formula uses a relativistic modification of d_{ca} . Once the cell width in QuickPIC shrinks below $\Delta_0/4$, the stopping power levels off. We did not see this behavior with the Bohr stopping power in Section 4.7.1 for the electrostatic case and where $n_0 \lambda_{De}^3$ is small. The saturation of the stopping power as $\Delta/\Delta_0 \rightarrow 0$ is worthy of future investigation. If the saturation is real, then Vlasov theory breaks down at scales before QED or small impact parameters become important.

As a final sanity check of the stopping power in QuickPIC, we increased the charge of the projectile for the largest cell-width. Doubling the charge multiplies the measured stopping power by a factor of four, and quadrupling the charge multiplies it by a factor of 16. This quadratic scaling of the stopping power with charge is what we expect based on the stopping power formulas, and demonstrates

that the plasma is responding linearly. If we increase the charge enough, the plasma will no longer respond linearly, and the scaling will no longer be quadratic. However, we will not need to concern ourselves with non-linear stopping power response except with the most densely packed electron beam configurations.

Since QuickPIC has proven to produce well-behaved stopping powers, it is well-suited for our correlated stopping simulations. While it does not allow us to study the effects of beam divergence and angular scattering, we can still study basic correlations effects in detail. We use QuickPIC for all our correlated stopping simulations in Chapter 6 except for one instance.

4.8 Conclusion

In this Chapter, we have investigated the single-particle stopping power of particles with speeds greatly in excess of the electron thermal velocity. This includes the stopping power of relativistic electrons. Various derivations are reviewed and compared, including the Bohr, Bethe, and QED expressions. We examine the stopping power predicted from the wake of a particle moving through a plasma represented by a Vlasov fluid and discuss how to include discrete particle or collisional effects at small impact parameter or large k (in k -space integrals). We also derive the wake and the stopping power associated with it for a relativistic electron moving in a cold plasma fluid. We use a derivation used in the plasma-based accelerator community.

We then conduct electrostatic simulations using ddcMD and BEPS for plasmas with very small plasma parameters ($n_0\lambda_{De}^3 \sim 1$). For the BEPS (a PIC code) simulations, we vary the particle size, including cases where it is much less than a Debye length. For ions, the ddcMD cases agree with the Bohr formula, and the BEPS results agree with the force from the wake on the particle for finite-size particles (this eliminates divergences). The ddcMD simulations show large jumps

in energy loss when close encounters occur. These jumps are not seen in BEPS, since the cell size, although small compared to a Debye length, is too large for collisions with large energy transfer. The BEPS results do converge to the ddcMD results as the cell size is reduced.

For electrons, the agreement with theory is not as good, because the energy transfer is very large from collisions with small impact parameters. These interactions “heat” the test charges, complicating the statistics. Using test charges that do not respond to the fields from the plasma could have yield cleaner results. For relativistic electrons, we need to use codes that have full EM fields and relativistic pushers. We tried OSIRIS, but the finite-difference time-domain field solvers support EM waves with phase velocities less than c on the grid, leading to unphysical Čerenkov radiation. We therefore use QuickPIC, a quasi-static code described in Appendix A. We find that QuickPIC provides stopping powers below those of the QED, Bohr, and Bethe expressions (as expected). As the particle size was reduced, the stopping power from QuickPIC (solely from the wake) saturated at values below the other expressions. Further investigations are needed to determine if this saturation is physical, i.e., whether discrete particle effects (not from collisions) come into play before quantum effects are important.

Future work includes examining how the conclusions change regarding the wake and collisional effects as the plasma parameter is increased. This work would include studying how energy is left behind the particle. It would be interesting to examine the relative amount of energy in collective or average motion versus random motion. This may be a better way of comparing the effects of the wake and collisions. It will also be of interest to use even smaller cells in QuickPIC as well as using improved field solvers in OSIRIS. Another possibility would be to use a spectral EM PIC code.

CHAPTER 5

The Anatomy of Relativistic Wakes

5.1 Introduction

In Chapter 3, we examined non-relativistic electrostatic wakes using Vlasov theory and particle codes. As a particle approaches the speed of light, electromagnetic effects become important, which change the behavior of the wake. In this chapter, we examine relativistic wakes (wakes with $v_{phase} \sim c$) through simulations, including the magnetic field contribution, and make comparisons with the fluid wake formula we derived in Section 4.3. We use the code QuickPIC, which is described in Appendix A. We will see that there are significant differences between relativistic wakes and electrostatic wakes, but magnetic fields are negligible for wakes made by a single electron. We will find that fluid theory describes off-axis behavior of the wake well for the first wavelength, but deviation from linear fluid theory occurs very quickly.

5.2 Longitudinal and Focusing Fields

In Section 4.3, we derived the formula for the E_z field of the wake of a relativistic point particle moving in the \hat{z} direction,

$$E_z = -2 \frac{q_t \omega_{pe}^2}{c^2} K_0 \left(\frac{\omega_{pe}}{c} r \right) \cos \left[\omega_{pe} \left(t - \frac{z}{c} \right) \right] \Theta \left(t - \frac{z}{c} \right). \quad (5.1)$$

The stopping power from one electron arises from E_z in Equation 5.1 evaluated at the location of the electron. However, to set the stage for studying correlated stopping, we now would like to know the focusing field on a particle moving in the \hat{z} direction with a speed near c , and the evolution of both the accelerating and focusing fields in the wake. Let $W_{\parallel} = E_z$ and $W_{\perp} = E_r - B_{\theta}$ be the longitudinal and focusing fields, respectively. Then, because the wake is a function of $ct - z$, the Panofsky-Wenzel theorem [PW56] says that

$$\frac{\partial W_{\parallel}}{\partial r} = \frac{\partial W_{\perp}}{\partial z}. \quad (5.2)$$

Stated another way, the Panofsky-Wenzel theorem says that W_{\parallel} and W_{\perp} can be obtained from a single potential,

$$\psi = -2 \frac{q_t \omega_{pe}}{c} K_0 \left(\frac{\omega_{pe}}{c} r \right) \sin \left[\omega_{pe} \left(t - \frac{z}{c} \right) \right] \Theta \left(t - \frac{z}{c} \right). \quad (5.3)$$

Therefore,

$$W_{\perp} = -\frac{\partial \psi}{\partial r} = -2 \frac{q_t \omega_{pe}^2}{c^2} K_1 \left(\frac{\omega_{pe}}{c} r \right) \sin \left[\omega_{pe} \left(t - \frac{z}{c} \right) \right] \Theta \left(t - \frac{z}{c} \right). \quad (5.4)$$

Just like E_z , the focusing field diverges as $r \rightarrow 0$, falls off quickly in r , and is perfectly periodic in $t - z/c$ behind the particle.

In PIC simulations, we do not use point particles. Rather, we use finite-size particles. In most cases, a PIC particle represents many real particles, so they are referred to as macro- or super-particles. Here, each simulation electron represents a single real electron but has a finite size, as was the case in the simulations in Section 4.7. To get the wake from an electron with a finite size, we need to integrate Equation 5.1, 5.3, or 5.4 (the Green's function) over the particle shape. For example,

$$\psi_S(r, \theta, \xi) = -2 \frac{q_t \omega_{pe}}{c} \times \int_{-\infty}^{\xi} d\xi' \int_0^{\infty} r' dr' \int_0^{2\pi} d\theta' K_0 \left(\frac{\omega_{pe}}{c} |\vec{r} - \vec{r}'| \right) \sin \left[\frac{\omega_{pe}}{c} (\xi - \xi') \right] S(r', \xi'), \quad (5.5)$$

where $\xi \equiv z - ct$. If particle is much smaller than a skin depth, then the electric field of a particle with a Gaussian shape on-axis scales as $q \frac{\omega_{pe}^2}{c^2} \ln \left(\frac{1.12c}{\sigma \omega_{pe}} \right)$, where σ is the standard deviation or particle size [MTE14].

Our understanding of wake dynamics in a cold fluid agrees rather well with QuickPIC simulations within the first wavelength of the wake. Figure 5.1 shows this agreement in plots of E_z and the focusing field. This simulation uses the same parameters as in Section 4.7.3, except the box has dimensions of $20.7575 \frac{c}{\omega_{pe}} \times 20.7575 \frac{c}{\omega_{pe}} \times 41.5150 \frac{c}{\omega_{pe}}$ and $512 \times 512 \times 1024$ cells, giving a cell width $\Delta = \Delta_0 = 0.0405 \frac{c}{\omega_{pe}}$. The QuickPIC results will saturate as r approaches a cell width due to the use of finite-size particles. However, as shown in Figure 4.11, the QuickPIC result saturates as $\Delta/\Delta_0 \rightarrow 0$. Therefore, physical effects can saturate the wake on-axis before finite-size particle effects come into play.

For small amplitude wakes, the focusing field of an electron moving near the speed of light is $-\partial\psi/\partial r = E_r - B_\theta$. However, E_r and B_θ make very unequal contributions to the focusing field. Figure 5.2 demonstrates that B_θ is over four orders of magnitude less than E_r , except for the slice at the particle's position in z . The self-fields of the particle, E_r and B_θ , are very large in this slice, but $E_r - B_\theta \sim 0$, so the self-field on the particle nearly vanishes at all values of r . These results suggest that the contribution of B_θ to the focusing field for particles inside the wake is negligible. We also note that, although B_θ is large near the test charge, it has no effect on cold plasma particles that start at rest. However, for a $10keV$ plasma, B_θ could have a non-negligible effect on the background plasma. The large E_r field of the test charge pushes out the plasma particles radially, creating the wake. However, if one transformed to a center of mass frame where the plasma electrons were drifting towards the test charge, then the magnetic fields of both the test charge and plasma electrons would need to be considered. For a plasma with finite temperature, full electromagnetic and not just Coulomb interactions are therefore necessary to study collisions.




Figure 5.1: The E_z field (left) and focusing field (right) of an electron wake in a QuickPIC simulation. The top row shows 2D slices of the fields, and the green dot represents the electron's position. Lineouts are taken along the black lines and plotted as red curves in the bottom row, whereas the fluid theory results are also plotted as dashed blue curves.




Figure 5.2: The focusing fields in the QuickPIC simulation. 2D slices of E_r and B_θ are on the left and right, respectively. The green dots represent the electron's position.

We have demonstrated that fluid theory accurately describes the wake, except it deviates from QuickPIC near the axis due to the finite-size particle effects and/or discrete particle effects. However, the agreement holds only for the first wavelength after the particle. In the next section, we will see that the nice uniform oscillation predicted by the linear fluid formulation quickly breaks down after the first wavelength.

5.3 Curvature and Spreading

Contrary to the predictions of linear fluid theory, the relativistic wake in PIC simulations spreads out with increasing ξ (distance behind the particle) and acquires an upward concavity in the longitudinal direction, as demonstrated in Figure 5.3. We might expect this transverse spreading because the electrostatic wakes in Chapter 3 displayed it, but they also displayed a downward concavity in the longitudinal direction. This discrepancy is due to thermal effects. Here, we explore why the upward curvature occurs and how thermal effects change it.

Understanding why the curvature occurs requires examining the trajectories




Figure 5.3: 2D slices of E_z , the focusing field, and the density fluctuation on the left, center, and right, respectively. The wake acquires an upward curvature as we look farther behind the particle. The green dots represent the electron's position.

of the plasma particles. We shift now to OSIRIS, which has the ability to track plasma particles. The OSIRIS simulation has the same parameters as the QuickPIC simulation, but uses the fourth-order solver discussed in Section 4.7.2, and a box that is $20.7575 \frac{c}{\omega_{pe}}$ with 512 cells on all sides, giving a cell-width Δ_0 . The OSIRIS and QuickPIC wakes agree well, as we can see in Figure 5.4. The wavelength from fluid theory does not agree with the simulations along the axis, which is likely due to nonlinearities we will discuss next. The wavelength agrees better as one moves away from the axis. In addition, the amplitude on-axis agrees reasonably well with the location of the particle, and the amplitude of the wake slowly decreases behind the particle, in contrast with theory. Overall, the agreement between theory, QuickPIC, and OSIRIS is very good, allowing us to shift to OSIRIS to study the wake wavefront curvature.

To get hints for the reasons for the discrepancies, we examine the particle track from an OSIRIS simulation. The discrepancy is due to the breakdown of linear fluid theory. So, either nonlinear or non-fluid-like effects associated with




Figure 5.4: The E_z field of a wake in the OSIRIS simulation is on the left. The blue spot at the front of the wake is the electron’s location. On the right is a comparison of the E_z field lineouts of the wake from QuickPIC in dashed-dotted green, from OSIRIS in dashed red, and from fluid theory in solid blue (for a point particle but off-axis by $\Delta_0/2$). The lineouts are taken along the electron’s trajectory.

discrete particles are at play. As we noted earlier, the amplitude of the wake at the location of the particle saturates as $\Delta/\Delta_0 \rightarrow 0$. Note that as Δ is reduced, the size of the background plasma electrons also gets smaller. Furthermore, in both the OSIRIS and QuickPIC simulations, the ion background is a fixed but discrete background. The plasma particle tracks from the OSIRIS simulation are plotted in the $z - ct$ frame of the beam in Figure 5.5. The plasma particles move very little in z and move predominantly in x . Therefore, the variable $z - ct$ can be viewed as time. The plasma particles have an initial average separation of $0.0405 \frac{c}{\omega_{pe}}$ and their initial positions in x can be judged using the axis at the top of the plot, but their displacements in x are $\sim 10^{-4} - 10^{-5} \frac{c}{\omega_{pe}}$. Therefore, we have shifted the tracks close together for ease of viewing the tracks of so many particles at once, and the scale for the displacement is on the bottom. The tracks show us that, when the test electron passes through the plasma, it causes the nearest plasma electrons to undergo a “nonlinear” oscillation. We put “nonlinear” in quotes to

make it clear that we mean oscillations not at the linear wavelength but that may not depend on the amplitude of the oscillation. The time for these particles to restore to their original positions is greater than half a linear oscillation period (they oscillate with a period $> 2\pi/\omega_{pe}$). Measuring the first half-wavelength of the oscillation for the two center tracks in Figure 5.5 gives us $\approx 4.5\frac{c}{\omega_{pe}}$. These two tracks do not have the greatest oscillation amplitude because they are within a cell width of the path of the beam electron. As we examine tracks further to the right or left on the plot, the first half-wavelength of oscillation approaches $\pi\frac{c}{\omega_{pe}}$. This early nonlinear motion of the particles closest to the path of the beam electron produces the extended density depletion right after the beam electron in density fluctuation plot in Figure 5.3. Therefore, the early nonlinear motion explains why the wake curvature starts. We will discuss and eliminate possible causes for the nonlinear oscillations.

Looking at the motion of the particles during second and third wavelengths of oscillation in Figure 5.5, we can see that the nonlinear motion begins to propagate outward. During the second wavelength, the amplitude of oscillation of the tracks out to the fourth track from the center die down. At the same time, the oscillations of the fifth and sixth tracks from the center have gained amplitude, and their peaks occur earlier than the peaks of tracks closer to the center, which were delayed due to their nonlinear behavior during the first wavelength. During the next oscillation, the amplitudes of the fifth and sixth tracks have died down, but those of the eight and ninth have increased, and their peaks again occur earlier than those further in. In this way, the energy of the original nonlinear motion propagates outward, causing the wake spreading and curvature. However, the outside tracks oscillate with the linear wavelength with a constant amplitude

One possibility is that the anharmonic motion is described by electron rings in radial oscillations, as explained by Dawson in his classic 1959 paper [Daw59]. The radial oscillation description is reasonable for the high phase velocity of the wake.




Figure 5.5: Tracks of plasma particles in the wake in an OSIRIS simulation with $\Delta = \Delta_0$. The tracks are plotted in the frame of the moving window, and are in the same plane as the slices in Figure 5.3. The tracks have been shifted close to each other for ease of viewing. The x -axis is at the top of the box, to indicate the initial x -positions of the particles before they are affected by the test charge. The Displacement* axis at the bottom of the plot should be used as a scale for measuring a particle's displacement from its initial position at the top of the plot.

We determine the applicability of Dawson's solution through a direct comparison. According to Dawson, radial oscillations are described by the equation

$$m_e \frac{d^2 R}{dt^2} = -\frac{2\pi n_0 e^2}{r_0 + R} [(r_0 + R)^2 - r_0^2], \quad (5.6)$$

where r_0 is the equilibrium radial position of the plasma electron and R is its displacement from equilibrium. We compare with the second track from the right of center of Figure 5.5. For that track, in Figure 5.3, we plot R/r_0 as a function of t and compare with a numerical result using Equation 5.6. The numerical result starts at the first peak in the oscillation Figure 5.5, around $z - ct = 16 \frac{c}{\omega_{pe}}$. This comparison demonstrates that the Dawson description clearly does not describe the nonlinear oscillations in the wake. Therefore, another explanation needs to be found. Another possible explanation is related to discreteness effects. The Dawson ring model asserts that electron rings do not cross and that the force on any ring is due solely to the amount of ions crossed. Furthermore, the amount of ion charge is calculated assuming the ions are a perfectly smooth fluid.

As the particle size Δ gets smaller than the interparticle spacing and the displacement is small, then force on inner electrons will deviate from the ring model for fluid ions. This deviation is related to electron-electron and electron-ion collisions. For the outer electrons, there are enough ions inside of them that they still appear to be uniformly distributed and the uniform fluid description is still reasonable. As we make $\Delta/\Delta_0 \rightarrow 0$, we have observed that the inner oscillations of the wake are at a higher frequency and the outer oscillations are at the linear frequency (wavelength). Quantifying and explaining these observations are areas for future work. To date, we have noticed that, as $\Delta/\Delta_0 \rightarrow 0$, the curvature gets smaller but the wavelength gets much shorter than the linear wavelength. However, increasing the number of plasma particles per cell to one in the simulations with the smallest cell-width, $\Delta = \Delta_0/8$, (while maintaining ω_{pe}) returns the wavelength to that calculated using fluid theory while still eliminating the curvature. This result demonstrates that the upward wake curvature is a likely




Figure 5.6: For the second track to the right of center in Figure 5.5, starting at the peak of the oscillation around $z - ct = 16\frac{c}{\omega_{pe}}$, the simulation track is in blue, while the numerical integration of Equation 5.6 is in red.

due to a discreteness effect.

We also believe that the curvature is due to the discreteness effects and the initial amplitude of displacements as a function of radius. We are left with the question of why we instead saw a downward (opposite) curvature in our electrostatic PIC simulations in Chapter 3. The simple answer is plasma temperature, which leads to temperature-dependent diffraction [Fah10]. In Figure 5.7, we see that the upward curvature disappears then turns downward as we increase the initial background plasma temperature to $10eV$, $100eV$, and $1keV$ ($v_{th} = 0.00442c$, $0.0140c$, and $0.0442c$, respectively). The curvature is almost completely gone when $T_e = 100eV$, and has reversed when $T_e = 1keV$.

We note that this wake curvature effect seen in Figure 5.3 may only occur in cold, classical plasmas. In our simulations, the upward concavity disappears by $T = 100eV$. In an electron plasma with $n_0 = 10^{26}$, the Fermi energy [Kar07]




Figure 5.7: 2D slices of E_z of a wake in OSIRIS for 10eV, 100eV, and 1keV plasma temperatures on the left, center, and right, respectively. The upward curvature turns downward as the plasma heats up.

$E_f = 790eV$, where

$$E_f = \frac{\hbar^2}{2m_e} (3\pi^2 n_0)^{3/2}. \quad (5.7)$$

By the time the plasma enters the classical regime, $T \gg E_f$, the concavity has reversed. Studies of other densities, along with a quantum mechanical study are necessary to determine if this effect will occur in other regimes.

To summarize, we have seen that the nonlinear or not linear motion of plasma particles causes the upward concavity of the particle wake during the first oscillation. This nonlinear motion transfers radially outward during successive oscillations, causing the wake to spread. Increasing the plasma temperature, and thus the group velocity, causes the curvature to disappear, and the curvature turns downward, like we saw in Chapter 3, when the temperature is high enough. We also believe the deviation from linear oscillations may be due to discreteness effects. As we make Δ/Δ_0 (particle size to interparticle spacing) very small, we see the oscillations near the axis are much faster than linear oscillations seen for the outer electrons. Understanding this behavior is an are for future work. In the next chapter, we will finally see how the wake causes correlated stopping to differ from single-particle stopping.

CHAPTER 6

Correlated Electron Stopping in Fast Ignition Plasmas

6.1 Introduction

Electron beam stopping power plays a critical role in Fast Ignition (FI) [THG94, AV09, ASB07, STL12], and can be the determining factor in whether or not a target ignites. In FI, lasers are used to compress a spherical pellet of deuterium and tritium fuel, as in central hot-spot inertial confinement fusion. However, the compression stage in FI only creates the densities required for ignition, and does not heat the core enough to start fusion burn. After compression, a high-intensity short-pulse laser is fired at a point on the target (or into a gold cone), as in Figure 6.1. The overdense laser-plasma interactions need to produce a beam of energetic (3-5 MeV) electrons, which travel into the deuterium-tritium core of the target. As these electrons slow down inside the dense core, they heat it, initiating thermo-nuclear burn. Under the right conditions, once the fusion process starts, it will lead to a self-propagating burn wave, which we call “ignition.” Assuming a properly assembled and compressed target, the primary determining factor in whether or not a target ignites is the energy that the electron beam deposits into a local region of the core. Because the electron beam stopping power is a critical factor in this energy deposition process, it is also a critical factor in the feasibility of the FI concept.

Current calculations and fluid simulations of FI use the single-electron stopping




Figure 6.1: The Fast Ignition concept. After the compression stage, a high-intensity short-pulse laser hits the pellet of the compressed fuel, creating a beam of energetic electrons. As the electrons slow down in the core, they heat it, starting the ignition process.

power given by the QED stopping formula, derived in Section 4.6, for each electron in the beam. Such studies predict that 1-3 MeV electrons are required. However, the lasers needed to generate the electron beam intensity required for ignition also will have high intensity and will therefore produce electrons with energies in the 10-30 MeV range. As we noted earlier, a part of the stopping power can be viewed as the energy lost by the particle as it excites the wake, or equivalently, the backward force on the electron from the electric field in the wake at the position of the particle. These aforementioned simulations ignore the fact that electrons can interact with the wakes made by earlier electrons. This interaction can affect the stopping power of electrons later in the beam that are inside the wake. We call this interaction and its effect on stopping power “correlated stopping.” The effect has been studied extensively in the field of ion beam stopping [Ari78, BA95, BA96, Ari00, MWS07], where it is known as the “vicinage effect,” but it has been

largely ignored by those studying electron stopping, aside from References [BD08, DF99], which ignore beam electron dynamics. Some groups have studied stopping power enhancements in the Fast Ignition scenario due to fluid instabilities [MF02]. However, these studies are performed using a fluid description, which cannot be applied to electron beams in the core of current Fast Ignition target designs due to the “low” beam density. In this Chapter, we present results demonstrating that processes similar to these fluid instabilities occur outside the regime where a fluid description applies. We emphasize that, in our studies, the inter-particle spacing of the beam electrons can be much larger or on the order of the wake wavelength and spot size, i.e. $\geq 1c/\omega_{pe}$. In this regime, the fluid description is not accurate and the interaction between individual wakes is important. We show that this correlation effect is significant even outside the fluid regime, and almost always leads to an increase in stopping power in cold plasmas, with enhancements in excess of 100 in some circumstances.

The increased stopping power is significant because the short-pulse laser energy required for ignition is inversely proportional to the square of the stopping power under ideal circumstances [STL12]. This dependence can be seen by considering optimal parameters given by Atzeni *et al.* [ASB07]. While recent work by Tonge *et al.* [TMM09] and May *et al.* [MTE14] shows that this calculation is overly simplistic, the final formula does offer a guide to the relationship between electron stopping power and other parameters, such as short-pulse laser energy and intensity. The electrons are generated by a laser with pulse length t_{opt} and are incident on a hot-spot with penetration depth $\rho\delta z_{opt}$ and radius r_{opt} . The laser has a peak intensity I_{0L} and the electron beam has an intensity I_{0f} . We assume that the electron beam is perfectly collimated and has a spectrum described by ponderomotive scaling [WKT92], giving a temperature $T_p = (\sqrt{1 + a_0^2} - 1)m_e c^2$, where $a_0 = \frac{eE_{0L}}{m_e\omega_0 c}$ is the normalized vector potential of the laser. The beam electrons have an average energy $\langle E \rangle$ and lose E_{DT} each as they pass through the

hot-spot. Given these parameters, if the beam delivers E_{stop} to the hot-spot, then it must start out with

$$E_{fast} = \frac{E_{stop}^2}{\alpha I_{0S}}, \quad (6.1)$$

where $\alpha = \pi r_{opt}^2 t_{opt}$ and I_{0S} is a scaling factor given by

$$I_{0S} \lambda_0^2 = 13.7 \text{GW} \left[\frac{I_{0f}}{I_{0L}} \frac{E_{DT}}{m_e c^2} \frac{T_p}{\langle E \rangle} \right]^2. \quad (6.2)$$

Given that E_{DT} appears in the denominator of Equation 6.1 and is directly proportional to the electron stopping power, we can clearly see that the required beam energy is inversely proportional to the square of the stopping power and directly proportional to the square of the average energy of the beam electrons. Stated another way, less energy is needed if the stopping power increases, and more energy is needed if each electron has too much energy.

6.2 Simulation Codes, Parameters, and Simplifying Assumptions

We perform the correlated stopping simulations using QuickPIC, which is a PIC code that uses a quasi-static approximation. The approximation assumes that the beam evolves on a time-scale much slower than the time required for the beam to pass by a plasma particle. The main requirement is that the beam particles all move near the speed of light in the same direction. For details about QuickPIC and the quasi-static approximation, see Appendix A.

QuickPIC proved to be the best code for our simulations because it has a Galilean moving window at c and includes non-radiative electromagnetic fields, but does not include radiative fields, so it is not subject to numerical Čerenkov radiation. The moving window causes the box to move with the beam, allowing for simulations of beam propagation over large time- and space-scales. The field solver includes some retardation effects associated with causality. The fields in

front of the relativistically moving test charge vanish. However, the field solver does not include radiation effects. There are some similarities to non-radiative Darwin solvers [KR71, NL76, BL85]. The ddcMD code used in Chapter 3 is not appropriate because it is electrostatic and thus does not include field retardation effects needed for causality (the fields extend in front of the particle). Finally, numerical Čerenkov radiation, discussed in Section 4.7.2, prevented us from using finite-difference EM PIC codes at the time. Numerical Čerenkov radiation was the final impetus for moving to QuickPIC, which does not suffer from the effect. Given these three requirements, QuickPIC proved to be the best tool available, and we use it for all simulations in this chapter except for one instance, where we use OSIRIS. We note that there has been much recent work on both controlling the numerical Čerenkov instability as well as minimizing numerical Čerenkov radiation of single electrons [LLT13]. Codes such as OSIRIS can now be used to extend the results in this Chapter.

The assumption that the beam electrons are all moving in the same direction is one of many simplifications we have made in our study. This uniform direction of travel is an important point because electron beams in FI usually have a significant divergence [RTT06, SSA04]. Other simplifications include ignoring all quantum mechanical effects, the effect of background plasma temperature, binary encounters with plasma particles on distances smaller than the cell size, and beam temperature. We briefly touch on possible consequences of non-zero plasma temperature in Section 6.3.

The baseline parameters for the correlated stopping simulations are in Table 6.1. The plasma electron density is chosen to match the density in the core of a FI target with a core mass density of 450g/cm^3 . As in our earlier simulations, for both the beam and plasma particles, a single particle in QuickPIC represents a single physical particle, with the caveat that QuickPIC uses a 2D plasma slice. We use these parameters in all correlated stopping simulations in this chapter unless

QuickPIC Simulation Parameters	
n_0	10^{26} cm^{-3}
T_e	0 (cold)
Box Dimensions	$41.5150 \frac{c}{\omega_{pe}} \times 41.5150 \frac{c}{\omega_{pe}} \times 83.03 \frac{c}{\omega_{pe}}$
Box Cells	$1024 \times 1024 \times 2048$ for $\Delta_0 = 0.0405 \frac{c}{\omega_{pe}}$
dt	$10\omega_{pe}^{-1}$
Beam Electron \vec{P}_0	$10m_e c \hat{z}$
Transverse Boundaries	Conducting, Reflecting

Table 6.1: Baseline parameters for the QuickPIC correlated stopping simulations.

stated otherwise.

6.3 Betatron Motion

For relativistic electrons, the speed of each electron and the phase velocity of their wakes are all near the speed of light. In addition, due to their large mass, the ions do not move for many oscillations within each wake. As a result, we can use the simulated single-electron wake at one instance in time to study the dynamics of another electron inside the wake by assuming that the phase velocity of the wake does not change.

To study the electron motion, we developed a MATLAB® [MAT12] script that uses the fields from a QuickPIC simulation to calculate the motion of a beam of electrons inside a wake. This script allows us to study the trajectories of many separate electrons without performing an equal number of QuickPIC simulations. In this section, we use the QuickPIC wake from Section 5.2.

We compare the motion of a beam electron in the wake as calculated by QuickPIC and MATLAB in Figure 6.2. The tracks are plotted over the wake’s focusing field. The beam electron under examination starts with $\vec{P}_0 = 10m_e c \hat{z}$, and the sim-




Figure 6.2: Betatron motion of an electron in a wake as calculated using QuickPIC (yellow line) and a MATLAB script (black line). The electron starts at the blue dot and its tracks are plotted over the focusing field.

ulations are run for $3500/\omega_{pe}$. Curiously, the QuickPIC track has some “jumps” in it, which we do not currently understand. While the agreement between the script and QuickPIC is not perfect, it is good enough for rough studies of electron motion. The discrepancy (excluding jumps) is due to the difference between the y axis being $z - ct$ or $z - v_0 t$, where v_0 is the initial speed of the particle.

To observe the motion of beam electrons in different parts of the wake, we perform five runs using the MATLAB script with a beam electron starting in a different position each time. We also note that, in reality, the wake is also moving relative to the particle, so the fields are not “static,” which is assumed when calculating tracks using the MATLAB script. We again start each electron with $\vec{P}_0 = 10m_e c \hat{z}$, and run the script for $20,000/\omega_{pe}$ each time. The result, with the tracks imposed over both the longitudinal and focusing fields, is plotted in Figure 6.3. As we expect based on examining the focusing field, some electrons move into the wake (focused) and oscillate while others are pushed out of the wake (defocused). Most interesting is the electron that starts lowest on the plot, which




Figure 6.3: Betatron motion of five independent electrons in a QuickPIC wake as calculated using the MATLAB script. The tracks are imposed over the longitudinal field on the left and over the focusing field on the right. The bright green dots near the top of the plots represent the position of the electron producing the wake.

begins to move out of the wake, but is pushed up by E_z (a negative E_z accelerates the electron), and then turns back in to the left, transitioning from a defocusing phase to a focusing phase. The motion we see in these tracks is the reason why we need to consider electron dynamics in our correlated stopping studies.

The electrons that become trapped in the wake oscillate both transversely and longitudinally. The transverse oscillations are referred to as betatron oscillations in particle accelerators. We use the MATLAB script to calculate the motion of a beam electron trapped in the wake along with the change in F_z (δF_z) due to the trapping wake over an oscillation cycle, until it returns to near its relative starting point in z . The results are plotted in Figure 6.4. While δF_z due to the wake is asymmetric about zero during the cycle, due to $|\min(\delta F_z)| > |\max(\delta F_z)|$, δF_z time-averages to zero. This fact must be true because the trapped electron finishes the cycle very near to its starting position. This zero time-average force results from the static nature of the wake, i.e., the total force on the particle in the




Figure 6.4: The path of a beam electron during a complete oscillation cycle, returning to its starting point, inside the wake produced by another beam electron, is on the left. The change in F_z during the cycle experienced by the electron due to its presence inside the wake is plotted on the right.

frame of the wake depends on a single conservative potential ψ . Therefore, any increase in stopping power results from the non-static character of the wake of an electron beam with multiple particles due to correlation effects, or the dynamical change in the wake from a background plasma, or nonlinear frequency shifts.

Our study of betatron motion has used a cold plasma so far, so we briefly consider the effects of background temperature on the motion. We know from our studies in Chapter 3 that the amplitude of background fluctuations can exceed the amplitude of the wake. Therefore, we can expect that, above a certain temperature, the wake of a single beam electron will not be enough to trap another beam electron.

We examine the effect of temperature on the transverse and longitudinal separation between two beam electrons, with one trapped in the wake of another, using OSIRIS. The simulation parameters are the same as those we used in the QuickPIC simulation, with the differences being a fourth-order finite-difference field solver and $dt = 0.02/\omega_{pe}$. While unphysical numerical Čerenkov radiation precludes a study of stopping power using this solver, as discussed in Section

4.7.2, it does not prevent a study of the wake behind the particle. The wakes we see at $t = 100/\omega_{pe}$ in the simulation with a 1keV background temperature are shown in Figure 6.5 along with the background E_z field, where we have used the subtraction technique described in Section 3.4 to view the wake. The transverse and longitudinal separations between the two beam particles are plotted in Figure 6.6 for four different background temperatures. Clearly, the trailing beam electron is trapped when the plasma is cold. The wake has less of an effect as we raise the temperature to just 10eV, and the trailing electron is almost completely de-trapped at 100eV. At 1keV, we can see that the wake has no effect. If we had decreased the initial transverse separation of the particles, we expect that the lack of trapping would have occurred at a higher temperature. Oddly, the longitudinal separation oscillates as we increase the temperature, which is likely due to the longitudinal waves in the background E_z field, which are caused by the particle boundary conditions of the moving window. Despite this issue, these simulations clearly show that temperature plays an important role, and any correlated stopping effects at higher temperatures will require the interaction of wakes from several beam electrons. This will be an area for future work.

6.4 Beam Scaling Parameter

In the correlated stopping simulations, we will see that the electron beam density is a determining factor in the effect of correlated stopping. Given the length scale of the wake, c/ω_{pe} , it is natural to similarly scale the beam density. We define the number of beam electrons in a cubic skin depth, $N_b = n_b(c/\omega_{pe})^3$, where n_b is the beam density in unscaled units. N_b for a range of beam and plasma densities is plotted in Figure 6.7, where, for understanding the results of Fast Ignition scenarios, the densities are scaled to the laser critical density, n_c , and N_b is scaled to the number of plasma electrons in a cubic skin depth at the laser




Figure 6.5: The E_z fields at $t = 100/\omega_{pe}$ for the 1keV OSIRIS simulation used to study beam electron motion in a thermal plasma. The subtracted field is on the left, showing the wakes, with the blue dots at the start of the wakes representing the beam electron positions. The unsubtracted field is on the right.




Figure 6.6: The motion of one beam electron in another's wake in OSIRIS simulations with various background temperatures. The transverse separation vs. time is on the left and the longitudinal separation on the right. The (blue, green, red, black) line corresponds to a simulation with a background temperature of (0eV, 10eV, 100eV, 1keV).




Figure 6.7: The number of beam electrons in a cubic skin depth for a range of beam and plasma densities. For reference, the red ticks on the x -axis are the plasma electron densities for a 450g/cm^3 core with $1\mu\text{m}$ (1ω) and $0.5\mu\text{m}$ (2ω) lasers. The (linestyle, N_b , laser) triplets are (solid white, 1, 1ω), (dash-dotted red, 64, 1ω), (dashed green, 1, 2ω), and (dotted black, 64, 2ω).

critical density, N_c .

6.5 Relativistic Fluid Streaming Instabilities in 1D

Although the fluid equations cannot accurately describe the physics in many of our simulations, streaming instabilities derived from them do help explain two trends we observe in Sections 6.8 and 6.9. In particular, we observe two distinct stopping power growth rates at different times in some simulations, and we also observe fluid-like instabilities occurring more rapidly with distance from the front of the beam. The fluid description of streaming instabilities explains both of these observations.

We consider two fluid-type streaming instabilities. One is due to longitudinal bunching and the other transverse focusing. We derive the longitudinal bunching streaming instability in this Section and cover transverse focusing in Section 6.6. We start with the linearized fluid equations along with Poisson's equation for the beam and plasma electrons. The fluid equations for the plasma are

$$\frac{\partial \delta v_p}{\partial t} = -\frac{eE_z}{m_e} \quad (6.3)$$

$$\frac{\partial \delta n_p}{\partial t} + n_0 \frac{\partial \delta v_p}{\partial z} = 0, \quad (6.4)$$

where δn_p and δv_p are the plasma density and velocity perturbations, respectively. The fluid equations for the beam must be written relativistically. The linearized continuity equation for the beam requires no special treatment, and is

$$\frac{\partial \delta n_b}{\partial t} + n_b \frac{\partial \delta v_b}{\partial z} + c \frac{\partial \delta n_b}{\partial z} = 0, \quad (6.5)$$

where n_b is the initial beam density, and δn_b and δv_b are the beam density and velocity perturbations, respectively. The momentum equation must be calculated relativistically before linearization. For the 1D instabilities where $\vec{v}_b = v_b \hat{z}$, we have

$$\begin{aligned} \frac{\partial p_b}{\partial t} + v_b \frac{\partial p_b}{\partial z} &= -eE_z \\ \frac{\partial \gamma_b v_b}{\partial t} + v_b \frac{\partial \gamma_b v_b}{\partial z} &= -\frac{eE_z}{m_e} \\ \gamma_b^3 \left(\frac{\partial v_b}{\partial t} + v_b \frac{\partial v_b}{\partial z} \right) &= -\frac{eE_z}{m_e} \\ \left(\frac{\partial v_b}{\partial t} + v_b \frac{\partial v_b}{\partial z} \right) &= -\frac{eE_z}{\gamma_b^3 m_e}, \end{aligned} \quad (6.6)$$

where p_b is the beam momentum and γ_b is the beam Lorentz factor. Linearizing and assuming that $v_b \approx c$ results in

$$\left(\frac{\partial \delta v_b}{\partial t} + c \frac{\partial \delta v_b}{\partial z} \right) = -\frac{eE_z}{\gamma_b^3 m_e}. \quad (6.7)$$

The evolution of the electric field is described by Ampere's Law and Gauss's Law.

Assuming overall charge neutrality,

$$\frac{\partial E_z}{\partial t} = -4\pi en_0 \delta v_p + 4\pi en_b \delta v_b + 4\pi en_b v_b \quad (6.8)$$

$$\frac{\partial E_z}{\partial z} = -4\pi e(\delta n_p + \delta n_b). \quad (6.9)$$

We can put the equations we have derived so far into a more useful form by combining them. We take the convective derivative of Equation 6.5 and substitute in Equation 6.7 followed by Equation 6.9.

$$\begin{aligned} \left(\frac{\partial}{\partial t} + c\frac{\partial}{\partial z}\right)^2 \delta n_b &= -n_b \frac{\partial}{\partial z} \left(\frac{\partial}{\partial t} + c\frac{\partial}{\partial z}\right) \delta v_b \\ \left(\frac{\partial}{\partial t} + c\frac{\partial}{\partial z}\right)^2 \delta n_b &= \frac{n_b e}{\gamma_b^3 m_e} \frac{\partial E_z}{\partial z} \\ \left(\frac{\partial}{\partial t} + c\frac{\partial}{\partial z}\right)^2 \delta n_b &= -\frac{4\pi n_b e^2}{\gamma_b^3 m_e} (\delta n_p + \delta n_b) \\ \left(\frac{\partial}{\partial t} + c\frac{\partial}{\partial z}\right)^2 \delta n_b + \frac{\omega_{pb}^2}{\gamma_b^3} \delta n_b &= -\frac{\omega_{pb}^2}{\gamma_b^3} \delta n_p, \end{aligned} \quad (6.10)$$

where ω_{pb} is the beam plasma frequency. We then take the time derivative of Equation 6.4 and substitute in Equation 6.3 followed by Equation 6.9.

$$\begin{aligned} \frac{\partial^2 \delta n_p}{\partial t^2} &= -n_0 \frac{\partial}{\partial z} \frac{\partial \delta v_p}{\partial t} \\ \frac{\partial^2 \delta n_p}{\partial t^2} &= \frac{en_0}{m_e} \frac{\partial E_z}{\partial z} \\ \frac{\partial^2 \delta n_p}{\partial t^2} &= -\frac{4\pi e^2 n_0}{m_e} (\delta n_p + \delta n_b) \\ \frac{\partial^2 \delta n_p}{\partial t^2} + \omega_{pe}^2 \delta n_p &= -\omega_{pe}^2 \delta n_b. \end{aligned} \quad (6.11)$$

We now make the same change of variables as we use in Sections A.2 and A.3 for the quasi-static approximation in QuickPIC. We define $s \equiv z$ and $\xi \equiv ct - z$. The derivative transforms are

$$\frac{\partial}{\partial z} = \frac{\partial}{\partial s} - \frac{\partial}{\partial \xi} \quad (6.12)$$

and

$$\frac{\partial}{\partial t} = c \frac{\partial}{\partial \xi}. \quad (6.13)$$

Using this transformation, Equations 6.10 and 6.11 become

$$\left(\frac{\partial^2}{\partial s^2} + \frac{k_{pb}^2}{\gamma_b^3}\right) \delta n_b = -\frac{k_{pb}^2}{\gamma_b^3} \delta n_p \quad (6.14)$$

and

$$\left(\frac{\partial^2}{\partial \xi^2} + k_{pe}^2\right) \delta n_p = -k_{pe}^2 \delta n_b, \quad (6.15)$$

respectively, where $k_{pj} \equiv \omega_{pj}/c$. These are two coupled harmonic oscillator equations. The beam oscillates in s and the plasma in ξ .

Using these equations, we consider the evolution of the beam due to longitudinal bunching under two conditions, the long-pulse, short propagation distance limit, and the short-pulse, long propagation distance limit. Analogous equations have been studied extensively in laser-plasma interactions [Mor97]. Under the long-pulse limit, we assume that the fast variation occurs in ξ (the oscillations in ξ remain resonant), so $\delta n = \bar{n}(\xi, s)e^{ik_{pe}\xi}$. With this enveloping, Equation 6.15 becomes

$$\begin{aligned} \left(\frac{\partial^2}{\partial \xi^2} + k_{pe}^2\right) \bar{n}_p e^{ik_{pe}\xi} &= -k_{pe}^2 \bar{n}_b e^{ik_{pe}\xi} \\ \left(\frac{\partial^2 \bar{n}_p}{\partial \xi^2} + 2ik_{pe} \frac{\partial \bar{n}_p}{\partial \xi} - k_{pe}^2 \bar{n}_p + k_{pe}^2 \bar{n}_p\right) e^{ik_{pe}\xi} &= -k_{pe}^2 \bar{n}_b e^{ik_{pe}\xi} \\ 2ik_{pe} \frac{\partial \bar{n}_p}{\partial \xi} &\approx -k_{pe}^2 \bar{n}_b, \end{aligned} \quad (6.16)$$

where we have assumed that $|\frac{\partial \bar{n}_p}{\partial \xi}| \ll k_{pe} \bar{n}_p$. Equation 6.14 becomes

$$\frac{\partial^2 \bar{n}_b}{\partial s^2} = -\frac{k_{pb}^2}{\gamma_b^3} \bar{n}_p, \quad (6.17)$$

where we have assumed that $|\frac{\partial \bar{n}_b}{\partial s}| \gg \frac{k_{pb}}{\gamma_b^{3/2}} \bar{n}_b$ (the oscillations in s are strongly driven and non-resonant). Taking $\partial^2/\partial s^2$ of Equation 6.16 and substituting in Equation 6.17, our final differential equation becomes

$$\frac{\partial^2}{\partial s^2} \frac{\partial \bar{n}_p}{\partial \xi} = -\frac{i}{2} \frac{k_{pe} k_{pb}^2}{\gamma_b^3} \bar{n}_p. \quad (6.18)$$

Equation 6.18 does not have a closed form solution, but we can derive an asymptotic solution. We first make a Fourier transform in ξ and define $\Gamma = \frac{k_{pe} k_{pb}^2}{\gamma_b^3}$,

so that

$$\frac{\partial^2}{\partial s^2} i k_\xi \tilde{n}_p = -i \frac{\Gamma}{2} \tilde{n}_p. \quad (6.19)$$

Equation 6.19 is just a simple harmonic oscillator. The solution is

$$\tilde{n}_p(s) = A \exp\left(i\sqrt{\frac{\Gamma}{2k_\xi}}s\right) + B \exp\left(-i\sqrt{\frac{\Gamma}{2k_\xi}}s\right). \quad (6.20)$$

We continue forward with the second term. Performing the inverse Fourier transform,

$$\bar{n}_p \sim \int_{-\infty}^{\infty} dk_\xi \exp\left(-i\sqrt{\frac{\Gamma}{2k_\xi}}s\right) \exp(ik_\xi\xi). \quad (6.21)$$

To approximate this integral, we could use saddle point integration (method of steepest descents) [AWH13, Mor97]. However, the leading term in the asymptotic expansion can be obtained quickly by using instead the stationary phase approximation. Under this approximation, we Taylor expand the phase around the stationary point, and pull the first term out of the out of the integral. Letting $\phi_{pha}(k_\xi) = -\sqrt{\frac{\Gamma}{2k_\xi}}s + k_\xi\xi$, the phase is stationary when

$$\frac{\partial\phi_{pha}(k_\xi)}{\partial k_\xi} = 0 = \frac{1}{2}\sqrt{\frac{\Gamma}{2}}\frac{s}{k_\xi^{3/2}} + \xi, \quad (6.22)$$

or

$$\begin{aligned} \frac{1}{\sqrt{k_{0\xi}}} &= \left(\frac{1}{2} + i\frac{\sqrt{3}}{2}\right) \left(\frac{2\xi}{s}\sqrt{\frac{2}{\Gamma}}\right)^{1/3} \\ k_{0\xi} &= \left(-\frac{1}{2} + i\frac{\sqrt{3}}{2}\right) \left(\frac{s}{2\xi}\sqrt{\frac{\Gamma}{2}}\right)^{2/3}, \end{aligned} \quad (6.23)$$

Where we have taken the appropriate cube root of -1. Plugging Equations 6.23 into 6.21 to obtain the first term of the expansion yields

$$\bar{n}_p \sim \exp\left[\left(\frac{\sqrt{3}}{4} - \frac{3i}{4}\right) \Gamma^{1/3} \xi^{1/3} s^{2/3}\right]. \quad (6.24)$$

Plugging in for Γ , the non-enveloped δn_p evolves asymptotically as

$$\delta n_p \sim \exp\left[\left(\frac{\sqrt{3}}{4} - \frac{3i}{4}\right) (k_{pe}\xi)^{1/3} \left(\frac{k_{pb}}{\gamma_b^{3/2}}s\right)^{2/3}\right] e^{ik_{pe}\xi}. \quad (6.25)$$

Equation 6.25 clearly exhibits the exponential growth in time (s) and also demonstrates that the instability growth rate increases with distance from the head of the beam, as we observe with the long beams in Section 6.9.

We now consider the short-pulse, long propagation limit. In this limit, we assume that the fast variation occurs in s (the s oscillations are resonant), so that $\delta n = \bar{n}(\xi, s)e^{ik_{pb}s/\gamma_b^{3/2}}$. With this new enveloping, we see ξ and s switch places in the derivation, along with k_{pe} and $k_{pb}/\gamma_b^{3/2}$. Equations 6.15 and 6.14 then become

$$\begin{aligned}\frac{\partial^2 \bar{n}_p}{\partial \xi^2} &\approx -k_{pe}^2 \bar{n}_b \\ 2i \frac{k_{pb}}{\gamma_b^{3/2}} \frac{\partial \bar{n}_b}{\partial s} &\approx -\frac{k_{pb}^2}{\gamma_b^3} \bar{n}_p,\end{aligned}\tag{6.26}$$

leading to the combined equation

$$\frac{\partial^2}{\partial \xi^2} \frac{\partial \bar{n}_p}{\partial s} = -\frac{i}{2} \frac{k_{pe}^2 k_{pb}}{\gamma_b^{3/2}} \bar{n}_p.\tag{6.27}$$

The math is almost identical to before, leading to

$$\delta n_p \sim \exp \left[\left(\frac{\sqrt{3}}{4} - \frac{3i}{4} \right) (k_{pe} \xi)^{2/3} \left(\frac{k_{pb}}{\gamma_b^{3/2}} s \right)^{1/3} \right] \exp \left(i \frac{k_{pb}}{\gamma_b^{3/2}} s \right).\tag{6.28}$$

This new equation exhibits the exponential growth that is present in the long beam limit, and again the growth increases with distance from the front of the beam. However, the exponents on $k_{pe}\xi$ and $\frac{k_{pb}}{\gamma_b^{3/2}}s$ have flipped. The overall growth rate for the beam is slower.

These two limits can be viewed as parts of a process that occurs for a single beam. Early in time, the beam propagates a short distance relative to its own length, with its evolution described by Equation 6.25. As the beam propagates further into the plasma, the distance it has covered eventually becomes much greater than the length of the beam, moving the beam into the short-pulse long propagation distance limit, with its evolution described by the slower growth rate of Equation 6.25. We see this transition occur in our simulations in Sections 6.8 and 6.9.

6.6 Transverse Beam Self-Modulation

When we examine the evolution of long beams in Section 6.9, we will see that the electrons in the focusing regions of the wake will bunch both transversely and longitudinally with time, while the electrons in the defocusing regions spread out. The time-evolution of the focusing and spreading leads to a spot-size modulational instability, which has been studied extensively in plasma-based accelerators. We use an analysis performed by Weiming An as part of his Ph.D. research [An13]. This analysis is for a single Gaussian beam. In FI, the beams are very wide, and a wide beam breaks up into many Gaussian beams. We briefly cover the primary results of his work in this Section.

We assume that the beam has a Gaussian transverse profile,

$$\rho_b = \frac{q_b \Lambda_b}{R^2} \exp\left(-\frac{x^2 + y^2}{2R^2}\right), \quad (6.29)$$

where q_b is the charge of a beam particle, Λ_b is the beam number density per unit length, and $R = \sqrt{\langle x^2 \rangle} = \sqrt{\langle y^2 \rangle}$ is the transverse spot size, where the wedge brackets indicate an average over all of the beam particles. The evolution of the spot-size is described by an envelope equation,

$$\frac{\partial^2 R}{\partial s^2} - \frac{q_b \langle x W_x \rangle}{\gamma_b m_b c^2 R} - \frac{\epsilon_N^2}{\gamma_b^2 k_{pe}^2 R^3} = 0, \quad (6.30)$$

where W_x is the transverse focusing field $\nabla_{\perp} \psi$, γ_b is the beam particle's Lorentz factor, and $\epsilon_N = \gamma_b k_{pe} \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle x x' \rangle^2}$ is the normalized beam emittance, where $x' = dx/ds$.

The equation for ψ can be derived from Equation 4.24. Using the fact that $E_z = -\partial\psi/\partial z$, we have

$$(\nabla_{\perp}^2 - k_{pe}^2)\psi = -4\pi e \delta n_p. \quad (6.31)$$

For wide beams, where $\nabla_{\perp}^2 \psi \ll k_{pe}^2 \psi$, it follows that

$$\delta n_p = \frac{k_{pe}^2}{4\pi e} \psi. \quad (6.32)$$

Plugging this Equation into Equation 6.15 yields

$$\left(\frac{\partial^2}{\partial \xi^2} + k_{pe}^2\right) \psi = -4\pi\rho_b. \quad (6.33)$$

To find the evolution of the envelope, we work with wide beams and in the long-beam limit, so that ψ will evolve adiabatically. Under this condition, we can perturb the beam about an equilibrium state, $\Lambda_b = \frac{m_b c^2 \epsilon_N^2}{\pi q_b^2 \gamma_b}$. We let $R = R_0 + R_1$ and $\psi = \psi_0 + \psi_1$, where 0 indicates an equilibrium and 1 indicates a perturbation, and $R_0 \gg R_1$ and $|\psi_0| \gg |\psi_1|$. Using $W_x = -\partial\psi_1/\partial x$, then,

$$\frac{\partial^2 R_1}{\partial s^2} + \frac{2\epsilon_N^2}{\gamma_b^2 k_{pe}^2 R_0^3} \frac{R_1}{R_0} = -\frac{q_b \langle x \partial\psi_1/\partial x \rangle}{\gamma_b m_b c^2 R_0} \quad (6.34)$$

and

$$\frac{\partial^2 \langle x \partial\psi_1/\partial x \rangle}{\partial \xi^2} + k_{pe}^2 \langle x \partial\psi_1/\partial x \rangle = -\frac{2m_b c^2 \epsilon_N^2}{q_b \gamma_b} \frac{R_1}{R_0^3}. \quad (6.35)$$

We define $M_1 = \frac{q_b \gamma_b k_{pe}^2 R_0^3 \langle x \partial\psi_1/\partial x \rangle}{2m_b c^2 \epsilon_N^2}$ and $k_s^2 = \frac{2\epsilon_N^2}{\gamma_b^2 k_{pe}^2 R_0^4}$, transforming Equations 6.34 and 6.35 into

$$\begin{aligned} \frac{\partial^2 M_1}{\partial \xi^2} + k_{pe}^2 M_1 &= -k_{pe}^2 R_1 \\ \frac{\partial^2 R_1}{\partial s^2} + k_s^2 R_1 &= -k_s^2 M_1. \end{aligned} \quad (6.36)$$

These equations have the same form as those in Section 6.5 and also are essentially the general form for streaming instabilities. We now assume that M_1 takes the form $M_1 = \bar{M}_1 e^{-ik_{pe}\xi} + c.c.$, where *c.c.* denotes the complex conjugate. Similarly, $R_1 = \bar{R}_1 e^{-ik_{pe}\xi} + c.c.$ We assume that $\partial M_1/\partial \xi \ll k_{pe} M_1$ and that $\partial R_1/\partial s \gg k_s R_1$. We can then write

$$\frac{\partial}{\partial \xi} \frac{\partial^2 \bar{R}_1}{\partial s^2} = \frac{i}{2} k_s^2 k_{pe} \bar{R}_1. \quad (6.37)$$

As in Section 6.5, we find an asymptotic solution to Equation 6.37 using the method of stationary phase,

$$R_1 = \bar{R}_1 e^{-ik_{pe}\xi} \sim \exp \left[\left(\frac{3\sqrt{3}}{4} + \frac{3i}{4} \right) (k_s s)^{2/3} (k_{pe} \xi)^{1/3} \right] e^{-ik_{pe}\xi}. \quad (6.38)$$

This Equation clearly tells us that the amplitude of the spot size modulations grow as the beam moves forward through the plasma and that the growth increases from head to tail, as we will see in Section 6.9.

6.7 A Beam of 125 Electrons with Various Densities

We next carry out QuickPIC simulations of correlated stopping. As a starting point in our study of correlated stopping, we examine 125 beam electrons inside a cubic space, all starting with $\vec{P}_0 = 10m_e c \hat{z}$. We perform multiple simulations, and vary N_b between 1/64 and 1,000. For reference, $N_0 \equiv n_0(c/\omega_{pe})^3 = 1.5 \times 10^4$ at $n_0 = 10^{26} \text{ cm}^{-3}$. Therefore, for the largest N_b we use, $n_b/n_0 = 4.3 \times 10^{-3}$. If $n_b = n_c$ for $1\mu\text{m}$ light, then $n_b/n_0 = 10^{-5}$ and $N_b = 0.13$. For each value of N_b , we perform a single run where all 125 beam electrons are on a lattice, and 16 where they are placed between zero and $\pm 1/2$ of the average electron separation from their lattice centers in each Cartesian direction. In all the simulations, we see the average stopping power increase with time except for the runs with $N_b = 1/27$ and $1/64$. We also note that, because we only used 125 beam electrons, the initial dimensions of the beam decrease from $20c/\omega_{pe}$ on each side when $N_b = 1/64$ to $0.5c/\omega_{pe}$ when $N_b = 1,000$.

6.7.1 One Beam Electron Per Cubic Skin Depth

The simple case of 125 beam electrons with $N_b = 1$ offers significant insight into electron beam dynamics. In particular, the simple lattice case illustrates both longitudinal and transverse bunching. We note that their effects can occur in the fluid limit, where they are commonly known as the beam-plasma instability [DMO70, GR71] and self-focusing [RSC89, RSC90], respectively, the former of which is related to the two-stream instability [Bun59, Sha63, OWM71], and the latter of which is related to the filamentation instability [TRT06, DHH72, LL73].

The filamentation instability operates through a mechanism similar to the Weibel instability [Wei59], and the two names are sometimes (incorrectly) used interchangeably. Past work has suggested that density gradients in Fast Ignition targets can be arranged to exploit these instabilities and decelerate ultrarelativistic electrons in the target core [MF02]. In our studies, particularly when $N_b \lesssim 1$, the fluid description is inaccurate and discrete interactions with wakes are important. However, we show that longitudinal and transverse bunching occur for $N_b \gtrsim 1/8$.

The simplest simulation we perform starts with the beam electrons on a $5 \times 5 \times 5$ lattice, as plotted in Figure 6.8. The forces at that time are in Figure 6.9 along with the E_z field, illustrating how the wakes interfere and affect the forces. The plot of the forces shows that some beam electrons will be accelerated while others are decelerated, bringing the electrons closer together longitudinally, which is the mechanism of the beam-plasma instability. Some forces are also pointing inward, and while these transverse forces oscillate due to betatron motion, they eventually cause filamentation. These features are present in most other beams we examine, although the beam self-focuses rather than filaments in these other simulations.

The evolution of the forces can be seen in Figure 6.10. By $t = 3,005/\omega_{pe}$, some of the longitudinal forces have greatly decreased or reversed, but filamentation isn't yet apparent. At $t = 6,895/\omega_{pe}$, we can see that the electrons have bunched together both longitudinally and transversely, resulting in much greater forces in $-\hat{z}$, which corresponds to a much greater stopping power, averaged over all electrons. The evolution of this stopping power is plotted in Figure 6.11, where it starts out less than the single-electron stopping power and shows a steady exponential growth with time.

The evolution of the stopping power changes when we randomly displace the initial positions of the beam electrons from their lattice centers, as plotted in Figure 6.12. The displacement in each Cartesian direction is a random number generated from a uniform distribution with a range of $[-1/2 \ 1/2)$ times the inter-




Figure 6.8: The 125 lattice beam electrons at $t = 5/\omega_{pe}$ with $N_b = 1$.




Figure 6.9: The relative forces on the 125 lattice beam electrons at $t = 5/\omega_{pe}$ with $N_b = 1$ are on the left. On the right, the forces are plotted over the E_z field in a slice across the middle of the box in y .




Figure 6.10: The relative forces on the 125 beam electrons for the lattice case with $N_b = 1$ are plotted on the left and right at $t = 3,005/\omega_{pe}$ and $6,895/\omega_{pe}$, respectively.




Figure 6.11: The stopping power for many different electron configurations, all with $N_b = 1$, averaged over all the beam electrons and plotted as a multiple of the single-electron stopping power. The horizontal dashed black line is the single-electron stopping power. The dashed purple line is the lattice case, and all other lines are for simulations where the electrons are initialized with random deviations from their lattice centers. The shaded areas represent the standard deviation for lines of the same color. On the right, there are 16 solid lines, corresponding to 16 different simulations.




Figure 6.12: The electron beam configuration at $t = 5/\omega_{pe}$ with the electrons displaced random distances from their lattice centers.

particle spacing. We use the native PythonTM [Pyt12] random number generator, which implements the Mersenne Twister algorithm [MN98] and is initialized with the operating system's randomness source. Using these random displacements, we see in Figure 6.11 that the stopping power starts near that of a single electron and grows very quickly. In these simulations, the electrons rapidly bunch transversely, and the initial growth in stopping power is faster than exponential. As the simulations evolve further, transverse and longitudinal spreading and re-bunching cause the stopping power to continue changing and to oscillate slightly. Eventually, the bunches spread apart transversely, starting at the tail of the beam, causing a steady decrease in stopping power.

A final factor worthy of investigation with this configuration is the cell width. We saw in Section 4.7.3 that changing the cell width (particle size) changes the stopping power for an individual electron. However, we do not expect the correlated stopping to depend on the particle size if many electrons are involved. We




Figure 6.13: A comparison of the correlated stopping power between $\Delta = \Delta_0$ and $\Delta_0/2$. with a decreased cell width. The stopping power is plotted as a multiple of the single-electron stopping power when $\Delta = \Delta_0/2$. The (dashed blue, dashed-dotted green) and (solid red, dotted magenta) lines are the lattice and random placement cases for a cell width of $(\Delta_0, \Delta_0/2)$, respectively. The shaded areas represent the standard deviation for lines of the same color. The horizontal (dash-dotted gray, dashed black) line is the single-electron stopping power for $\Delta = (\Delta_0, \Delta_0/2)$.

decrease the cell width by half to $\Delta_0/2$ and perform two simulations, one with the lattice case and one with a random placement case. These beam electron configurations are the two we use in the left plot in Figure 6.11. As we see in Figure 6.13, the smaller cell width does make a slight difference, but given the chaotic nature of the problem and the significant increase in computation requirements, it is not worth using in the majority of our studies.

The results of changing the cell width also indicate that the correlation effects are additive rather than multiplicative, i.e., it depends on $NeNe/Nm_e \sim Ne^2/m_e$.

Therefore, the single-electron stopping power does not matter much when studying correlated stopping. This result is not very surprising since correlated stopping is a wake effect, and wakes are largely independent of cell width except within a couple cell widths from $r = 0$, as we discussed in Section 5.2. We can also use this fact to bound the multiplicative effect of correlated stopping relative to classical stopping powers, which we do later.

The simulations with 125 beam electrons with $N_b = 1$ have demonstrated that correlation effects increase stopping power via discrete particle-like beam-plasma instabilities including filamentation or self-focusing and the beam-plasma-like instability. The growth is much more rapid when the electrons are placed a random distance from their lattice centers than in the lattice case. The faster growth is due to a larger noise source for the instabilities to grow from. Decreasing the cell width changes the evolution of the stopping power a little, but not enough to justify using a smaller cell width as we move forward with more complex simulations. This simple case has illustrated the mechanisms of correlated stopping, and we will now see how correlated stopping affects stopping power as we vary the density and size of the beam.

6.7.2 Various Beam Electron Densities

As we vary the beam electron density, we expect the stopping power to vary as well. If we start the 125 electrons closer together, the tighter bunch will result in more constructive interference of the wakes, and the natural longitudinal and transverse bunching will further enhance the stopping power. In addition, if we keep the number fixed but decrease the spacing such that $N_b \gg 1$, then all the beam electrons will be in the same phase of the wake as the first beam electron. In contrast, starting the electrons farther apart can prevent bunching from occurring at all, preventing any correlation effects except in the lattice case. We now examine simulations with various initial electron separations to observe the impact on




Figure 6.14: For various initial electron beam densities using 125 electrons, the average stopping power for the lattice cases is plotted on the left, and the stopping power averaged across all 16 random placement runs is plotted on the right. The (line style, N_b) pairs are (dash-dotted yellow, 1/64), (dashed black, 1/27), (solid blue, 1/8), (dotted green, 1), (dash-dotted red, 8), (dashed cyan, 64), (solid magenta, 1,000). The shaded areas represent the standard deviation for lines of the same color.

stopping power, and see that the maximum stopping power increases to 90 ± 50 times the single-electron stopping power for $N_b = 1,000$.

In Figure 6.14, for each initial beam electron separation, we plot the stopping power of the lattice case and the stopping power averaged over all 16 random placement cases. As the initial electron separation decreases, the stopping power peaks at a higher value and occurs earlier in time. The early stopping powers from the random placement and lattice cases also move into closer agreement. This convergence occurs because, when all the electrons are placed in the same phase of the wake, no noise source is needed before correlation effects are present.

One caveat is the transverse spreading of the beam, especially with the higher values of N_b . For large N_b and fixed beam electron number, all the electrons reside

in the decelerating part of the total wake, slowing down significant numbers in the beam. These decelerated beam electrons then phase slip to the defocusing region, causing the beam to blow apart. Some beam electrons hit the conducting wall and bounce off it. This reflection leads to the rapid wiggles in the stopping power of some of the lattice runs. In the next Section, we increase the number of beam electrons as N_b increases to reduce the artificial seeding of correlation effects. This change will prevent all of the electrons being bunched into the decelerating phase of the wake when N_b is large.

We provide bounds for the effect of correlated stopping in Figure 6.15 using the maximum measured stopping powers. We use the single-electron stopping power from QuickPIC to provide the upper bound since it lies below the results from all the stopping power formulas (we plot the ratio of stopping power to the single-particle result). The Bohr formula provides a lower bound since it provides the greatest single-electron stopping power.

We have observed that bunching of the beam electrons increases stopping power, and starting the beam electrons closer together leads to higher peak stopping power. The obvious question concerns how these effects translate to larger beams, which we consider next.

6.8 Various Beam Densities Initially Inside a 1,000 $(c/\omega_{pe})^3$ Cube

Similar to the increase in stopping power we observed by moving the 125 beam electrons closer together, the stopping power increases when we maintain the initial volume of the beam but increase the beam density. The effect is more drastic in this new situation because there are more electrons that can bunch together due to instabilities. In the simulations in this section, the initial electron beam is a cube with $10c/\omega_{pe}$ on each side. We vary N_b from 1/8 to 64. We do




Figure 6.15: For various initial electron beam densities using 125 electrons, the average peak stopping power is plotted as a multiple of the single-electron stopping power measured in QuickPIC for the lattice cases in dashed red and random placement cases in solid blue. The modified multiple of the Bohr stopping power is plotted for the lattice cases in dash-dotted magenta and for the random placement cases in dotted black.

not explore $N_b = 1,000$ due to numerical difficulties. The peak stopping powers jump significantly at the higher beam densities, going up to 400 ± 600 times the single-electron stopping power for the $N_b = 64$ random placement cases.

Beginning with these simulations, we take into account the additional complication of beam electrons leaving the box during the simulation. This factor makes the stopping power calculation marginally more difficult. All electrons are initially part of the calculation, but we explicitly discount them when they leave the box.

We plot the average stopping power over the beam for the lattice cases and the average over all 16 random placement cases for each beam density in Figure 6.16. The bounds on the maximum stopping power using the QuickPIC single-electron stopping power and the Bohr stopping power are presented in Figure 6.17. For the higher beam densities, the stopping power peaks at much higher values than when we only used 125 electrons, but it also falls off much faster late in time. The increased number of beam electrons leads to the possibility of more constructive interference between the wakes, and thus stronger fields overall.

The stopping powers in the cases of $N_b = 64$ and 8 exhibit rapid initial exponential growth, then slower exponential growth, and the random placement cases experience a plateau in between. This behavior is consistent with beam-plasma-like instabilities evolving from early time to late time growth, as discussed in Section 6.5. The first burst of growth is caused almost entirely by self-focusing or spot size self-modulation. There is then a bit of a lull before the longitudinal bunching beam-plasma-like instability takes over, causing the second burst of growth. This order is the opposite of what we observed with the lattice case in Section 6.7.1. The process is understandable, however, because the self-focusing aligns the electrons transversely, which increases the E_z amplitude, allowing the beam-plasma-like instability to take hold more easily.

To understand some simple scaling, we discuss purely temporal growth. The




Figure 6.16: For various electron beam densities initially inside $1,000 (c/\omega_{pe})^3$, the average stopping power for the lattice cases is plotted on the left, and the stopping power averaged across all 16 random placement runs is plotted on the right. The (line style, N_b) pairs are (solid blue, $1/8$), (dotted green, 1), (dash-dotted red, 8), and (dashed cyan, 64). The shaded areas represent the standard deviation for lines of the same color.




Figure 6.17: For various electron beam densities initially inside $1,000 (c/\omega_{pe})^3$, the average peak stopping power is plotted as a multiple of the single-electron stopping power measured in QuickPIC for the lattice cases in dashed red and random placement cases in solid blue. The modified multiple of the Bohr stopping power is plotted for the lattice cases in dash-dotted magenta and for the random placement cases in dotted black.

temporal growth rate of the beam-plasma instability is

$$\frac{\delta_{BP}}{\omega_{pe}} = \frac{\sqrt{3}}{2^{4/3}} \left(\frac{n_b}{n_0} \right)^{1/3} \frac{1}{\gamma_b}. \quad (6.39)$$

Later, we will compare the growth to the full spatio-temporal growth. Here, we compare the temporal growth rate to the measured growth rate of the stopping power (integrated over ξ) in the random placement simulations. We consider the second period of growth, which we believe is due to longitudinal effects. For $N_b = 64$, we measure $\frac{\delta_{BP}}{\omega_{pe}} = 4.08 \times 10^{-3}$ between $t = 290/\omega_{pe}$ and $560/\omega_{pe}$, less than half the 0.011 given by the formula. $N_b = 8$ exhibits greater disagreement, with a measured value of $\frac{\delta_{BP}}{\omega_{pe}} = 1.32 \times 10^{-3}$ between $t = 1200/\omega_{pe}$ and $1700/\omega_{pe}$, less than 1/5 of the formula's 7.3×10^{-3} . Unfortunately, the stopping power curve for the $N_b = 1$ case exhibits a downward concavity after $t \sim 3000/\omega_{pe}$, as seen in Figure 6.16, about when the beam-plasma-like instability appears to take hold. The instability is not apparent in the stopping power curve during the simulation time for $N_b = 1/8$. An important reason for this increasing discrepancy between simulation and theory as the beam density decreases is particle discreteness. As the beam density decreases, there are fewer beam electrons inside an electron wake wavelength and radius. While $N_b = 64$ is outside the fluid regime, the lower beam densities are further outside it, increasing the disagreement with the fluid growth rate. Other contributing factors include the small width of the beam and the fact that the tail of the beam is spreading, as shown in Figure 6.18 and discussed later.

Figure 6.18 clearly demonstrates self-focusing and the beam-plasma-like instability occurring in one of the random placement simulations with $N_b = 64$. The plots show electron positions and the beam density at $t = 600/\omega_{pe}$, the time of peak stopping power. The instabilities have led to beam densities far exceeding the initial density at a point where the electrons have bunched both longitudinally and transversely. For comparison, the initial average beam density is $4 \times 10^{-3}n_0$. These high beam densities lead to the impressive stopping powers we observe.

The E_z field and plasma return current (J_z) are plotted in Figure 6.19 at the time of peak stopping power for the same simulation as in Figure 6.18. The upward wake concavity that we discussed in Section 5.3 is present but not obvious. There is also no observable damping of the wake or the current, which is expected. The wake $v_{phase} \sim c$, so there are no resonant electrons, and hence no Landau damping [Lan46]. The current has a slight negative bias, so that current neutrality is maintained.

The phase space density as a function of z and P_z is plotted at $t = 600/\omega_{pe}$ in Figure 6.20 for the same simulation as in Figure 6.18. The distribution clearly shows the result of the beam-plasma-like instability, with a curve starting at the front of the beam with $P_z = 10m_e c$ but decreasing further back in z . Then, after $z = 673c/\omega_{pe}$, where we see the big bunch in the density plot, there is a spread in momentum, with $P_z > 10m_e c$ for many electrons. This distribution is the result of electrons in the head of the bunch slowing down while those in its tail speed up, causing the bunching.

The plots in Figure 6.18 show the beam exploding transversely in addition to bunching. The larger initial beam densities lead to strong focusing and de-focusing fields, so parts of the beam pinch rapidly, but other parts explode quickly. As more of the beam spreads out, the stopping power falls off.

This spreading also means that electrons begin hitting the conducting wall early in the simulation. At the time of peak stopping power, beam electrons have already hit the conducting wall in the $N_b = 64$ case, as clearly seen in Figure 6.18. However, when only a few electrons reach the wall, the average stopping power is not significantly affected when thousands of electrons are in the beam. Once the beam has diffused enough for the wall interaction to become important, the simulation has progressed enough for the stopping power to have decreased far below the peak.

The simulations of various beam densities initially inside a $1,000 (c/\omega_{pe})^3$ cube




Figure 6.18: The electron positions are plotted on the left in 3D and the beam density is plotted on the right for a 2D slice through the middle of the box in y . The plots are at $t = 600/\omega_{pe}$, and the run is one of the random electron placement cases with $N_b = 64$ initially inside the $1,000 (c/\omega_{pe})^3$ cube.




Figure 6.19: The E_z field and the \hat{z} component of the plasma current for a slice through the middle of the box in y are plotted on the left and right, respectively, at $t = 600/\omega_{pe}$ for the simulation in Figure 6.18.




Figure 6.20: The phase space density as a function of P_z and z at $t = 600/\omega_{pe}$ for the simulation in Figure 6.18.

have demonstrated that correlated stopping can increase the stopping power several hundred times above the single-electron stopping power at higher beam densities. They have also demonstrated that the growth in stopping power occurs first through self-focusing then through the beam-plasma-like instability. While useful, the beams are still small compared to a true electron beam. In the next Section, we extend the length of the beam, which further increases the stopping power.

6.9 Box-Length Beams

Electron beams in Fast Ignition can stretch for hundreds of microns to centimeters, and therefore exceed thousands of box lengths for our simulation parameters. In a move towards this limit, we examine the evolution of beams that stretch along the entire length of the box. We would also prefer the beam to cover the box transversely, but refrain from expanding it to minimize interactions with the conducting boundaries. Instead, we again use beam dimensions of $10c/\omega_{pe}$ in x and y .

Peak Stopping Powers and Time-Steps		
N_b	dt [ω_{pe}^{-1}]	Max dE/dx [single-electron dE/dx]
1/8	5	2 ± 2
1	2	6 ± 15
8	2	100 ± 300
64	2	$(1 \pm 3) \times 10^3$

Table 6.2: The peak stopping powers of a typical electron in box-length beams averaged over the eight random placement simulations for each value of N_b , along with the time-steps used in those simulations.

We again vary N_b between 1/8 and 64. Due to the stronger fields from the longer beams, we use smaller time-steps for these simulations, which are in Table 6.2 along with the peak stopping powers averaged over the eight random placement simulations. The stopping powers for the higher beam densities are significantly higher than those in Section 6.8. In contrast, the stopping power for $N_b = 1/8$ hardly changes (but may be limited by the relatively small transverse beam size).

We plot the average stopping power for the lattice cases and the average over all eight random placement cases for each electron beam density in Figure 6.21. The bounds on the maximum stopping power using the QuickPIC single-electron stopping power and the Bohr stopping power are in Figure 6.22. Extending the beam causes the stopping power to grow more rapidly and peak earlier in the simulation. The stopping power growth is again exponential and occurs in two phases, self-focusing followed by the beam-plasma-like instability, with the former having a faster growth rate.

We again compare the theoretical temporal growth rate of the beam-plasma instability with the second observed growth rate (of the stopping power) in the random placement simulations. We consider the second period of growth, which we believe is due to longitudinal effects. For $N_b = 64$, we measure $\frac{\delta_{BP}}{\omega_{pe}} = 0.035$

between $t = 80/\omega_{pe}$ and $110/\omega_{pe}$, over three times the formula's value of 0.011. For $N_b = 8$ we measure $\frac{\delta_{BP}}{\omega_{pe}} = 5.6 \times 10^{-3}$ between $t = 200/\omega_{pe}$ and $400/\omega_{pe}$, slightly less than the formula's 7.3×10^{-3} . As in Section 6.8, we are unable to find good time ranges to measure the growth rate of the beam-plasma-like instability for $N_b = 1$ and $1/8$. These measurements are a stark contrast to those in Section 6.8, where the measured temporal growth rates for $N_b = 64$ and 8 were significantly below the theoretical ones. The growth rate increases with distance from the front of the beam, causing the greatly increased temporal growth measured here. We discussed the theory of this spatio-temporal growth in Section 6.5 and explore it in the simulations in this Section.

There are secondary maximum in the stopping power curves, which occur at values of t when the peaks occurred for the shorter beams in Section 6.8. These secondaries suggest that, while tail of the beam is diffusing transversely after bunching, the beam-plasma-like instabilities are still taking place at the head of the beam. However, these secondary peaks are not as high as the peaks of the shorter beams, likely due to electrons further back in the beam experiencing a relatively lower stopping power. Examining E_z provides support for this hypothesis. Figure 6.23 shows the transverse average of E_z and the lineout down the middle of the box along z at the times when the stopping power peaks occur for a run with randomly placed beam electrons and $N_b = 64$. The first peak occurs at $t = 140/\omega_{pe}$ and the second at $t = 600/\omega_{pe}$. During the first peak, the field amplitude increases from the head of the beam to the tail. During the second peak, the E_z amplitude increases in the front half of the beam when compared to the field during the first peak, and decreases in the trailing half. The change is most dramatic near the center of the box in the transverse directions, where the field amplitude changes from increasing from head to tail to being relatively constant.

We examine the timings of the peaks in stopping power in Table 6.3. Let




Figure 6.21: For various electron beam densities in the box-length beams, the average stopping power for the lattice cases is plotted on the left, and the stopping power averaged across all eight random placement runs is plotted on the right. The (line style, N_b) pairs are (solid blue, $1/8$), (dotted green, 1), (dash-dotted red, 8), and (dashed cyan, 64). The shaded areas represent the standard deviation for lines of the same color.




Figure 6.22: For various electron beam densities inside box-length beams, the average peak stopping power is plotted as a multiple of the single-electron stopping power measured in QuickPIC for the lattice cases in dashed red and random placement cases in solid blue. The modified multiple of the Bohr stopping power is plotted for the lattice cases in dash-dotted magenta and for the random placement cases in dotted black.




Figure 6.23: The transverse average of E_z and the lineout down the middle of the box along z at the times when peak average stopping power occurs. The beam electrons are randomly placed and $N_b = 64$.

Times of Peaks			
N_b	Beam electron spacing [c/ω_{pe}]	Peak 1 [ω_{pe}^{-1}]	Peak 2 [ω_{pe}^{-1}]
1/8	2	1200	No Data
1	1	1150	No Data
8	0.5	460	2000
64	0.25	140	600

Table 6.3: The peak timing for the average stopping power from the random placement runs at each N_b . The secondary peaks do not appear during the simulations for $N_b = 1$ and $1/8$.

t_{p1-N_b} be the time of the first peak for a given N_b . Then, $t_{p1-1}-t_{p1-8} = 690/\omega_{pe}$ and $t_{p1-8}-t_{p1-64} = 320/\omega_{pe}$. The first difference is roughly twice the second. This doubling suggests that the timing of the first peaks depends inversely on the beam inter-particle spacing when $N_b \geq 1$, which is also the scaling of the beam-plasma instability growth rate. Unfortunately, we lack the data to examine the density dependence of the second peak timing.

The electron positions and density at the time of peak stopping power are in Figure 6.24 for the simulation examined in Figure 6.23. These plots show clear pinching and longitudinal bunching of the beam, which are evidence of self-focusing and the beam-plasma-like instability, respectively. The peaks in beam density occur relatively far back in the beam compared to the peak in Section 6.8, as expected in a spatio-temporally growing instability like that discussed in Section 6.5. The beam is also expanding radially, again with an increasing rate farther back in the beam, as expected from the transverse self-modulation discussed in Section 6.6.

The E_z field and plasma return current (J_z) are plotted in Figure 6.25 at the time of peak stopping power for the same simulation as in Figure 6.23. As we saw in Figure 6.23, the E_z amplitude increases from the head to the tail of the beam.




Figure 6.24: The electron positions are plotted on the left in 3D and the beam density is plotted on the right for a 2D slice through the middle of the box in y . The plots are at $t = 140/\omega_{pe}$ for the same run as in Figure 6.23.

The current amplitude also increases from the head to the tail of the beam. The beam electrons are clearly moving in the $+\hat{z}$ direction, so the plasma electrons are moving in the $-\hat{z}$ direction. The plasma electrons provide a return current to compensate for the incoming current.

There is clear evidence of the beam-plasma-like instability in the phase-space distribution in Figure 6.26, given the significant spread in momentum. The distribution is plotted at the time of peak stopping power for the simulation in Figure 6.23. The spread increases from the head of the beam to the tail, indicating that the instability grows from the head to the rear of the beam (it depends on ξ). This fact, combined with the conical transverse spreading of the beam in real space and the secondary peak in stopping power discussed earlier, suggest that the instabilities grow from the head to the tail and that the growth in ξ increases with s , as predicted by the spatio-temporal theory discussed in Section 6.5. For a beam several microns in length, the instability would rapidly saturate. Effects such as wave-breaking [KM88, AP56] and Langmuir wave collapse [Zak72, Ber98] may come into play. It may also be affected by a confinement mechanisms, such as strong magnetic fields, or a larger transverse size. Furthermore, the beams




Figure 6.25: The E_z field and the plasma z current in a slice through the middle of the box through y are plotted on the left and right, respectively, at $t = 140/\omega_{pe}$ for the same simulation as in Figure 6.23.

in FI are many microns wide. Therefore, filamentation rather than whole-beam self-focusing will be dominant.

We have demonstrated that beams covering the length of the box show a greatly increased stopping power compared to the shorter ones in Section 6.8. These longer beams have shown that the self-focusing and beam-plasma-like instability start strongest and peak first at the rear of the beam and move forward. The detailed operation of these instabilities in larger and longer beams, as well as beams with realistic and longitudinal momentum spreads, remains unclear, and is an area for future research.

6.10 Conclusion and Future Work

Our simulations clearly demonstrate that correlation effects on relativistic electron beam stopping should be considered. In the simulations of box-length beams, we observed the average stopping power increase to 1,000 times the single-electron stopping power for the beam with initial $N_b = 64$. The causes of this increase are discrete particle-like self-focusing and the beam-plasma instability. For $N_b \gtrsim 64$,




Figure 6.26: The phase space density as a function of P_z and z at $t = 140/\omega_{pe}$ for the same simulation as in Figure 6.24.

fluid-like descriptions for the beam-plasma interactions are reasonable. While such a beam density may be extreme for the core of an FI target, all our simulations indicate that there is an increase in stopping power for $N_b \geq 1/8$. Clearly, ignoring the coherent interactions of discrete particle wakes and the related self-focusing, filamentation, and beam-plasma-like instability leaves out important factors in the stopping power. In particular, because correlated stopping depends on N_b , stopping power may be great enough for FI to be feasible using cores with densities lower than $450\text{g}/\text{cm}^3$. This work indicates that understanding better the transition from the coherent interaction of discrete wakes to fluid-like streaming instabilities needs to be undertaken, including revisiting ideas such as those of Malkin and Fisch [MF02].

Clearly, these effects need further research. We have left out many factors, including background temperature, beam divergence, angular scattering, and energy spread, all of which will change the results. Necessary work that is straight-forward includes studying wider beams and adding a beam temperature (in all directions). A further improvement would be to use a periodic box, which would allow us to

expand the beam to fill the box transversely without concerns about the conducting boundary conditions. In addition, recent improvements in field solvers may soon permit using OSIRIS or other full PIC codes. These codes would allow for studies of beams with energies from a few MeV to 10s of MeV and large transverse momentum spreads.

CHAPTER 7

Conclusions and Future Work

PIC simulations have been used to study two processes involving wakes in plasmas, backward stimulated Raman scattering (BSRS) and electron beam stopping. In SRS, the beating of the pump and scatter light induces a plasma wave wake in the plasma which, in turn, scatters more light from the pump. In electron beam stopping, each electron produces a wake which contributes to its own stopping power, and the wakes of electrons earlier in the beam alter the stopping power of electrons later in the beam in a process dubbed “correlated stopping.” The studies produced many new and interesting results with important implications for Inertial Confinement Fusion (ICF).

The BSRS studies of Chapter 2 demonstrated that the convective amplification of a scattered light seed pulse in 1D OSIRIS simulations does agree with linear gain calculations. For the first time, special relativity was accounted for in the linear gain calculations, producing excellent agreement with the OSIRIS simulations when a relativistic particle pusher was used. The relativistic gain rates exceed the non-relativistic rates and peak at shorter wavelengths. This result has important implications for envelope codes such as pF3D, as it implies that they produce results with the wrong wavelength and amplitude in the linear regime.

However, the method used in this dissertation to account for special relativity in the gain calculation involves numerically integrating the Jüttner-Synge distribution over all values of the Lorentz factor for each point on the gain curve. The tested approximations to the relativistic dielectric produced unacceptable results.

While the integration method proved acceptably fast for calculating a single gain curve, a much better algorithm will be needed to account for relativistic gain in envelope codes used for hohlraum simulations.

After the seed exits the box in the BSRS simulations, the wake induced in the plasma by the seed and the pump continues to grow and traps particles, leading to a reduction in Landau damping and causing kinetic inflation with its accompanying large reflectivities later in the simulation. Kinetic inflation does not occur without a seed. This result suggests a process by which kinetic inflation may occur outside of laser speckles. Scattered light or a plasma wave produced via BSRS involving a hot-spot are the most obvious potential instigators. If the scattered light or plasma waves travel outside the speckle or remain after the speckle passes, they can initiate kinetic inflation in lower-intensity parts of the beam. These kinetic inflation results should be studied in more detail outside of 1D collisionless PIC simulations, as discussed in Section 2.9.

Extending the duration of the seed in the BSRS simulations resulted in kinetic inflation while the seed is still present as well as negative amplification of the seed when the nonlinear phase shift between the wake and scatter light is sufficiently large. For seeds far from resonance, the reflectivity amplitude oscillates with a frequency equal to the difference between the seed and resonant frequencies.

The studies of plasma wakes left behind by individual particles in Chapters 3, 4, and 5 demonstrate excellent agreement between the electrostatic wake in 3D BEPS and Vlasov theory. In relativistic simulations, the cold fluid theory accurately describes the wavelength and the general shape of the wake within the first wavelength, but misses the transverse spreading and upward concavity of the wake. These deviations from fluid theory are likely due to discrete particle effects rather than nonlinear physics. This concavity requires further study, as the simulations were classical, while a quantum mechanical description is more appropriate for a cold plasma with densities of 10^{26}cm^{-3} . The upward concavity

also disappears and turns downward as the plasma temperature increases, and identifying the reason for this change is left for future study.

In Chapter 6, 3D PIC simulations demonstrated for the first time that correlation effects have a significant impact on electron beam stopping power in Fast Ignition (FI) plasmas. Filamentation and two-stream instabilities (in both discrete particle and fluid-like regimes) dynamically increase the stopping power depending on the size and density of the beam. When beam electrons are initialized with random placement, filamentation occurs first, followed by the two-stream instability, resulting in the stopping power growing at two visibly distinct rates. In long electron beams, with dimensions $10c/\omega_{pe} \times 10c/\omega_{pe} \times 80c/\omega_{pe}$, the time-scale of the instabilities depends on the distance from the front of the beam, so that they occur most rapidly at the tail of the beam. The beam eventually spreads out transversely, leading the stopping power to drop off. For a long beam with an inter-particle separation of $0.25c/\omega_{pe}$, the peak stopping power of a typical beam electron is $(1 \pm 3) \times 10^3$ times that of an uncorrelated electron.

The correlated electron stopping results are a starting point for further research. The simplest studies that should be performed in the future involve increasing the beam size and studying the effects of beam temperature. The effects of background temperature, angular scattering, and beam divergence all need to be examined. Studies of the propagation of realistic FI beam distributions must be performed. Once a better understanding of correlated stopping is developed, applying the work of Malkin and Fisch [MF02] to exploit the instabilities may be possible.

This dissertation improves the understanding of the complex role that wakes play in plasmas, and in ICF in particular. The BSRS results have an important impact on the understanding of laser back-scatter in ICF plasmas, where the decrease in energy reaching the target hinders progress towards ignition. In contrast, wakes play a beneficial role in correlated electron stopping, which is a

largely unexplored path towards ignition via the FI concept. Both sets of results assist progress towards ignition in ICF.

APPENDIX A

QuickPIC

A.1 Introduction

QuickPIC [HDR06, ADM13] is a particle-in-cell (PIC) code specifically designed to simulate the evolution of ultra-relativistic charged particle beams in plasmas. It uses a quasi-static approximation, allowing it to take time-steps that are orders of magnitude larger than those of ordinary electromagnetic PIC codes. This approximation assumes that the time for the beam evolution is much greater than the time required for a beam particle to pass a plasma particle. Under this assumption, the time-step only needs to be small enough to resolve the evolution of the beam. In this Appendix, we review aspects of the quasi-static approximation and some details about the QuickPIC implementation without going into the details of the numerical implementation.

A.2 Basic Setup and Implementation Details

The key insight of the quasi-static approximation is that it decouples the beam push, i.e., advance from the calculation of the plasma response. We start with a 3D beam moving in the z -direction with a speed near c . At the beginning of a time-step, we deposit the beam charge and calculate the fields from it. Next, we calculate the plasma response using a 2D sheet of plasma particles. The sheet starts at the right side of the box, in front of the beam, and streams across the box to the left. As the sheet streams across the box, at each cell, we self-consistently

calculate the response of a plasma particle to the beam and other plasma particles in the sheet using a 2D solver. This plasma response gives us the wake, which we then use to push the beam forward. Another key point is that the charge on each plasma particle changes to account for plasma particles moving in the \hat{z} direction.

To make the quasi-static approximation, we change variables from the standard Cartesian coordinates (x, y, z, t) to (x, y, s, ξ) , where $s \equiv z$ and $\xi \equiv ct - z$. We can view the beam as evolving in s and the plasma as evolving in ξ . In the implementation, plasma particles move “forward” in ξ as they move across the box, as illustrated in Figure A.1 along with the corresponding plasma response. If a plasma particle has a very small value of V_z compared to c , then, in the “frame” of a beam moving at approximately c , the variable $\xi \approx ct$. Therefore, the variable ξ is time-like for the plasma particles. Conversely, the beam moves forward at c so ξ does not change much. However, $s \approx ct$ for the beam particles. This relationship allows us to move the box forward at the speed of light, commonly called a “moving window,” and treat s as a time-like variable for the beam. Whenever the beam takes a time-step, the box takes an s -step. Within each s -step, each ξ -step is like a 2D time-step for the plasma particles within a sheet as it crosses the box. The complete 3D cycle with the 2D cycle embedded inside it is shown in Figure A.2.

A.3 The Quasi-Static Approximation

With the information about the variables and implementation from Section A.2, we can understand what is happening as we examine the mathematical foundation behind the quasi-static approximation. We follow the work of Huang *et al.* [HDR06] and add some conceptual details.

As we have mentioned several times, the basic assumption behind the quasi-static approximation is that the beam evolves very slowly compared to the time it takes for the beam to pass a plasma particle. We can state this assumption




Figure A.1: A schematic of a 2D plasma sheet moving across the box in QuickPIC. As it crosses the box, it responds to the beam, producing the wake. The green dots are the beam electrons.




Figure A.2: The 3D flow chart of QuickPIC with the 2D flow chart embedded inside it. The 3D flow chart is repeated for each s -step. During each s -step, the 2D routine is called, and the 2D cycle repeats for each ξ -step as the 2D plasma sheet crosses the box.

mathematically using the variables s and ξ that we introduced in Section A.2. We can then write the assumption as

$$\frac{\partial}{\partial s} \ll \frac{\partial}{\partial \xi}. \quad (\text{A.1})$$

A more physical interpretation of the math is that the scale of the variation of quantities at a given part of the beam or wake (at a value of ξ) in s is much smaller than the wavelength of the wake.

Applying the chain rule and our assumptions, we can derive a couple of useful relationships between (s, ξ) and (z, t) .

$$\frac{\partial}{\partial z} = \frac{\partial}{\partial s} - \frac{\partial}{\partial \xi} \quad (\text{A.2})$$

and

$$\frac{\partial}{\partial t} = c \frac{\partial}{\partial \xi}. \quad (\text{A.3})$$

A.4 The Quantity ψ

At this point, it is useful to introduce the quantity ψ , which is commonly used in plasma accelerator physics. This single quantity relates the longitudinal field E_z and the focusing field $W_{\perp} = E_r - B_{\theta}$. This potential can be written as

$$\psi \equiv \phi - A_z, \quad (\text{A.4})$$

where A_z is the z component of the vector magnetic potential. ψ takes this form in any gauge. We find E_z by using Equations A.2 and A.3.

$$\begin{aligned} E_z &= -\frac{\partial \phi}{\partial z} - \frac{1}{c} \frac{\partial A_z}{\partial t} \\ &= -\frac{\partial \phi}{\partial s} + \frac{\partial \phi}{\partial \xi} - \frac{\partial A_z}{\partial \xi} \\ &= \frac{\partial \psi}{\partial \xi} - \frac{\partial \phi}{\partial s} \\ &\approx \frac{\partial \psi}{\partial \xi} \end{aligned} \quad (\text{A.5})$$

We can now use a more general version of the Panofsky-Wenzel theorem from Equation 5.2,

$$\nabla_{\perp} E_z = \frac{\partial \vec{W}_{\perp}}{\partial \xi}. \quad (\text{A.6})$$

This theorem allows us to find the focusing field by taking the perpendicular gradient of ψ ,

$$\vec{W}_{\perp} = -\nabla_{\perp} \psi = -\nabla_{\perp} \phi + \nabla_{\perp} A_z. \quad (\text{A.7})$$

It is also straight-forward to derive Equation A.7 from the Lorentz force equation by substituting for \vec{E} and \vec{B} in terms of the derivatives of the potentials and then using Equations A.2 and A.3. ψ appears regularly in the equations we derive in this Appendix.

A.5 Equations of Motion

Since the beam “moves” in s and the plasma particles “move” in ξ , we write separate equations of motion for the beam and plasma particles. We assume that the beam particles move at approximately c , which allows us to write $s \approx ct$ and $dt \approx ds/c$. For the beam’s transverse motion, we use Equation A.7 for the focusing field, which allows us to write,

$$\frac{d\vec{P}_{b\perp}}{ds} = -\frac{q_b}{c} \nabla_{\perp} \psi. \quad (\text{A.8})$$

For the position,

$$\frac{d\vec{x}_{b\perp}}{ds} = \frac{\vec{V}_{b\perp}}{c} = \frac{\vec{P}_{b\perp}}{\gamma_b m_e c}. \quad (\text{A.9})$$

Similarly, using Equation A.5 for the longitudinal field,

$$\frac{d\vec{P}_{bz}}{ds} = \frac{q_b}{c} \frac{\partial \psi}{\partial \xi}. \quad (\text{A.10})$$

The evolution of the ξ positions of the beam particles in the box are given by

$$\frac{\partial \xi_b}{\partial s} = \frac{1}{c} \frac{d}{dt} (ct - z) = 1 - \frac{1}{c} \frac{dz}{dt} = 1 - \frac{P_{bz}}{\gamma_b m_e c}. \quad (\text{A.11})$$

For the plasma particles, we hold s constant as we calculate the motion. The evolution of a particle's transverse momentum as it crosses the box in ξ is given by

$$\frac{d\vec{P}_{p\perp}}{d\xi} = \frac{dt}{d\xi} q_p \left[\vec{E}_\perp + \left(\frac{\vec{V}_p}{c} \times \vec{B} \right)_\perp \right] = \frac{q_p}{c - V_{pz}} \left[\vec{E}_\perp + \left(\frac{\vec{V}_p}{c} \times \vec{B} \right)_\perp \right]. \quad (\text{A.12})$$

The evolution for the transverse positions of the particles is described by

$$\frac{d\vec{x}_{p\perp}}{d\xi} = \frac{dt}{d\xi} \frac{d\vec{x}_{p\perp}}{dt} = \frac{\vec{V}_{p\perp}}{c - V_{pz}} = \frac{\vec{P}_{p\perp}}{\gamma_p m_e (c - V_{pz})}. \quad (\text{A.13})$$

The longitudinal momentum of the plasma particles is found from the constant of motion [MA97],

$$\gamma_p - \frac{P_{pz}}{m_e c} = 1 - \frac{q_p \psi}{m_e c^2}. \quad (\text{A.14})$$

If we use $\gamma_p^2 = 1 + \frac{\vec{P}_p^2}{m_e^2 c^2}$, then, after some algebra,

$$\frac{P_{pz}}{m_e c} = \frac{1 + \frac{\vec{P}_{p\perp}^2}{m_e^2 c^2} - \left(1 - \frac{q_p \psi}{m_e c^2}\right)^2}{2 \left(1 - \frac{q_p \psi}{m_e c^2}\right)}. \quad (\text{A.15})$$

In principle, a plasma particle has some slight s dependence, which is easily found by inverting Equation A.11 and modifying the notation appropriately,

$$\frac{ds_p}{d\xi} = \frac{1}{1 - \frac{P_{pz}}{\gamma_p m_e c}} = \frac{1}{1 - \frac{V_{pz}}{c}}. \quad (\text{A.16})$$

We assume that this s dependence is negligible in the equations of motion and don't integrate this equation. However, we do take it into account in the plasma charge and current deposition.

The implementation in QuickPIC casts the equations of motion into a “non-relativistic” form and uses a Boris pusher [Bor71]. This transformation involves modifying the fields and the charge. We only modify the transverse equations of motion, A.12 and A.13. We rewrite the momentum equation in terms of the quantity $\vec{u}_p \equiv \gamma_p \vec{V}_p$, which leads to

$$\frac{d\vec{u}_{p\perp}}{d\xi} = \frac{q_{eff}}{m_e c} \left[\gamma_p \vec{E}_\perp + \left(\frac{\vec{u}_p}{c} \times \vec{B} \right)_\perp \right], \quad (\text{A.17})$$

where the effective charge is defined as

$$q_{eff} \equiv \frac{q_p}{\gamma_p \left(1 - \frac{V_{pz}}{c}\right)} = \frac{q_p}{\left(\gamma_p - \frac{P_{pz}}{m_e c}\right)} = \frac{q_p}{1 - \frac{q_p \psi}{m_e c^2}}. \quad (\text{A.18})$$

For the transverse position, we rewrite the denominator in the same way as we just did for q_{eff} ,

$$\frac{d\vec{x}_{p\perp}}{d\xi} = \frac{\vec{u}_{p\perp}}{\gamma_p (c - V_{pz})} = \frac{\vec{u}_{p\perp}}{c \left(1 - \frac{q_p \psi}{m_e c^2}\right)}. \quad (\text{A.19})$$

An observant reader may notice an interesting dilemma with the plasma particle equations of motion. We have regularly stated that QuickPIC self-consistently calculates the response of the plasma to the beam. However, the equations of motion for the plasma particles depend on the fields, and when the plasma particles start streaming across the box, the only fields we have are from the beam. The response of the plasma to the beam fields will create new fields. The plasma's response should include those fields as well. When we say that the plasma response is found self-consistently, we mean that we have taken into account its own fields and those from the beam. QuickPIC resolves this problem by using a predictor-corrector method to calculate the plasma response, which it iterates for accuracy.

While the longitudinal momentum equation is not modified in the same manner as the transverse momentum equation, we re-write it to directly yield u_{pz} and γ_p . u_{pz} immediately follows from Equation A.15,

$$\frac{u_{pz}}{c} = \frac{1 + \frac{\vec{u}_{p\perp}^2}{c^2} - \left(1 - \frac{q_p \psi}{m_e c^2}\right)^2}{2 \left(1 - \frac{q_p \psi}{m_e c^2}\right)}. \quad (\text{A.20})$$

Plugging Equation A.15 into A.14 immediately yields γ_p ,

$$\gamma_p = \frac{1 + \frac{\vec{u}_{p\perp}^2}{m_e^2 c^2} + \left(1 - \frac{q_p \psi}{m_e c^2}\right)^2}{2 \left(1 - \frac{q_p \psi}{m_e c^2}\right)}. \quad (\text{A.21})$$

We now have the complete set of equations that describe the motion. We are missing the equations for the fields necessary to calculate the motion, which we derive next.

A.6 Field Equations

Like all PIC codes, QuickPIC deposits the charge and current, then uses them to solve for the fields. QuickPIC neglects the transverse beam current since it is small compared to the longitudinal beam current and unnecessary to for satisfying the quasi-static continuity equation.

The charge and current from the beam are deposited directly. However, since the plasma particles move in ξ rather than s , we must take this fact into account when depositing their charge and current. In Cartesian space, the charge density in a volume ΔVol is given by

$$\rho = \frac{1}{\Delta\text{Vol}} \sum_i q_i, \quad (\text{A.22})$$

where

$$\frac{1}{\Delta\text{Vol}} = \frac{1}{\Delta x \Delta y \Delta s}. \quad (\text{A.23})$$

When we deposit the plasma particles, we change from Δs to $\Delta\xi$ by using Equation A.16 and making the approximation

$$\frac{1}{\Delta\text{Vol}_p} = \frac{1}{\Delta x \Delta y \Delta\xi} = \frac{1}{\Delta x \Delta y \Delta s} \frac{\Delta s_p}{\Delta\xi} \approx \frac{1}{\Delta\text{Vol}} \frac{1}{1 - \frac{V_{pz}}{c}} \quad (\text{A.24})$$

Therefore, the plasma charge and current densities are given by,

$$\rho_p = \frac{1}{\Delta\text{Vol}} \sum_i \frac{q_{pi}}{1 - \frac{V_{pz}}{c}} = \frac{1}{\text{Volume}} \sum_i \gamma_{pi} q_{eff} \quad (\text{A.25})$$

and

$$\vec{J}_p = \frac{1}{\Delta\text{Vol}} \sum_i \frac{q_{pi} \vec{V}_{pi}}{1 - \frac{V_{pz}}{c}} = \frac{1}{\Delta\text{Vol}} \sum_i \vec{u}_{pi} q_{eff}, \quad (\text{A.26})$$

where we have used \vec{u}_p , γ_p , and q_{eff} from the end of Section A.5.

We mentioned in Section A.4 that ψ is used heavily in plasma accelerator physics, and we saw it appear frequently in the equations of motion. Therefore, our first step is to find the potentials so we can calculate it. We begin with the wave equations for the electrostatic potential ϕ and vector magnetic potential \vec{A} ,

$$\left(\frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \nabla^2\right) \phi(x, y, z, t) = 4\pi\rho(x, y, z, t), \quad (\text{A.27})$$

and

$$\left(\frac{1}{c^2} \frac{\partial^2}{\partial t^2} - \nabla^2\right) \vec{A}(x, y, z, t) = \frac{4\pi}{c} \vec{J}(x, y, z, t), \quad (\text{A.28})$$

where ρ and \vec{J} are the charge and current densities, respectively. Directly applying Equations A.2 and A.3, the wave equations become

$$-\nabla_{\perp}^2 \phi(x, y, z, t) = 4\pi\rho(x, y, z, t) \quad (\text{A.29})$$

and

$$-\nabla_{\perp}^2 \vec{A}(x, y, z, t) = \frac{4\pi}{c} \vec{J}(x, y, z, t). \quad (\text{A.30})$$

Armed with ϕ and \vec{A} , we can calculate ψ via Equation A.4. We can also relate ψ to \vec{A}_{\perp} via the Lorentz gauge condition,

$$\begin{aligned} \nabla \cdot \vec{A} + \frac{1}{c} \frac{\partial \phi}{\partial t} &= 0 \\ \nabla_{\perp} \cdot \vec{A}_{\perp} &= -\frac{1}{c} \frac{\partial \phi}{\partial t} - \frac{\partial A_z}{\partial z} \\ \nabla_{\perp} \cdot \vec{A}_{\perp} &= -\frac{\partial \phi}{\partial \xi} + \frac{\partial A_z}{\partial \xi} \\ \nabla_{\perp} \cdot \vec{A}_{\perp}(x, y, s, \xi) &= -\frac{\partial \psi(x, y, s, \xi)}{\partial \xi}. \end{aligned} \quad (\text{A.31})$$

ψ gives us E_z and W_{\perp} , but as explained in Section 5.2, W_{\perp} only applies to the beam particles. The electric fields are given by

$$E_z = \frac{\partial \psi}{\partial \xi} \quad (\text{A.32})$$

and

$$\vec{E}_{\perp} = -\nabla_{\perp} \phi - \frac{1}{c} \frac{\partial \vec{A}_{\perp}}{\partial t} = -\nabla_{\perp} \phi - \frac{\partial \vec{A}_{\perp}}{\partial \xi}. \quad (\text{A.33})$$

For the magnetic fields, we simply use the formula $\vec{B} = \nabla \times \vec{A}$ and transform the z derivative using Equation A.2.

$$B_z = [\nabla_{\perp} \cdot (\vec{A}_{\perp} \times \hat{z})] \hat{z} \quad (\text{A.34})$$

and

$$\vec{B}_{\perp} = \left(\nabla_{\perp} A_z - \frac{\partial \vec{A}_{\perp}}{\partial z} \right) \times \hat{z} = \left(\nabla_{\perp} A_z + \frac{\partial \vec{A}_{\perp}}{\partial \xi} \right) \times \hat{z}. \quad (\text{A.35})$$

An equation for $\frac{\partial \vec{A}_{\perp}}{\partial \xi}$ can be found by taking the ξ derivative of A.30,

$$-\nabla_{\perp}^2 \frac{\partial \vec{A}_{\perp}}{\partial \xi} = \frac{4\pi}{c} \frac{\partial \vec{J}_{\perp}}{\partial \xi}. \quad (\text{A.36})$$

Equation A.36 combined with A.17 helps clarify the issue of self-consistently finding the plasma response to the beam. The two equations together directly demonstrate the inter-dependence of the fields and the plasma particle motion. As we mentioned in Section A.5, QuickPIC resolves this issue by using a predictor-corrector method. Further details can be found in References [HDR06, ADM13].

A.7 Conclusion

By making the approximations we covered in this Appendix, QuickPIC is an excellent code for studying the evolution of ultra-relativistic electron beams. It can perform simulations using orders of magnitude less CPU time than fully electromagnetic PIC codes. However, it does have restrictions. We mentioned that it assumes that the beam is traveling in the \hat{z} direction, which means it is not useful for studying beam divergence and will have limited usefulness in studies of angular scattering.

In this dissertation, we use QuickPIC to examine the wake made by a single relativistic electron, i.e., the “beam” is one electron. We examine how the wake changes as the electron size is reduced and the spacing of the plasma electrons

increases. We also use it to study how a “beam” of discrete electrons interacts through the mutual interaction of their wakes.

Despite its limitations, QuickPIC has proved extremely useful in our studies due to its efficiency and because it resolves the issue with numerical Čerenkov radiation we discussed in Section 4.7.2. It has allowed us to perform the simulations in Chapter 6 within a very short time-span. These benefits led us to use it as our main tool to study correlated stopping.

REFERENCES

- [ADM13] Weiming An, Viktor K. Decyk, Warren B. Mori, and Thomas M. Antonsen Jr. “An improved iteration loop for the three dimensional quasi-static particle-in-cell algorithm: QuickPIC.” *Journal of Computational Physics*, **250**:165 – 177, 2013.
- [An13] Weiming An. *Efficient Modeling of Plasma Wake Field Acceleration Experiments Using Particle-In-Cell Methods*. PhD thesis, University of California, Los Angeles, 2013.
- [AP56] A. I. Akhiezer and R. V. Polovin. “Theory of wave motion of an electron plasma.” *Soviet Physics, Journal of Experimental and Theoretical Physics*, **3**(5):696–705, 1956.
- [Ari78] N. R. Arista. “Energy loss of correlated charges in an electron gas.” *Physical Review B*, **18**(1):1–8, Jul 1978.
- [Ari00] Néstor R. Arista. “Stopping of molecules and clusters.” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, **164–165**:108 – 138, 2000.
- [ASB07] Stefano Atzeni, Angelo Schiavi, and Claudio Bellei. “Targets for direct-drive fast ignition at total laser energy of 200-400kJ.” *Physics of Plasmas*, **14**(5):052702, 2007.
- [ASD09] S Atzeni, A Schiavi, and J R Davies. “Stopping and scattering of relativistic electron beams in dense plasmas and requirements for fast ignition.” *Plasma Physics and Controlled Fusion*, **51**(1):015016, 2009.
- [AV09] S. Atzeni and J. Meyer-ter Vehn. *The Physics of Inertial Fusion: Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter*. The International Series of Monographs on Physics. Oxford University Press, USA, 2009.
- [AWH13] George B. Arfken, Hans J. Weber, and Frank E. Harris. “Chapter 12 - Further Topics in Analysis.” In *Mathematical Methods for Physicists*, pp. 551 – 598. Academic Press, Boston, seventh edition, 2013.
- [BA95] E. M. Bringa and N. R. Arista. “Collective effects in the energy loss of ion beams in fusion plasmas.” *Physical Review E*, **52**(3):3010–3014, Sep 1995.
- [BA96] E. M. Bringa and N. R. Arista. “Energy loss of correlated ions in plasmas: Collective and individual contributions.” *Physical Review E*, **54**(4):4101–4111, Oct 1996.

- [BD08] Antoine Bret and Claude Deutsch. “Correlated stopping power of a chain of N charges.” *Journal of Plasma Physics*, **74**(05):595–599, 2008.
- [BE01] Jan Bergman and Bengt Eliasson. “Linear wave dispersion laws in unmagnetized relativistic plasma: Analytical and numerical results.” *Physics of Plasmas*, **8**(5):1482–1492, 2001.
- [Ber94] Jean-Pierre Berenger. “A perfectly matched layer for the absorption of electromagnetic waves.” *Journal of Computational Physics*, **114**(2):185–200, 1994.
- [Ber98] Luc Bergé. “Wave collapse in physics: principles and applications to light and plasma waves.” *Physics Reports*, **303**(56):259 – 370, 1998.
- [BHP08] Alfredo Bermdez, Luis HervellaNiето, Andrs Prieto, and Rodolfo Rodriguez. “Perfectly Matched Layers.” In Steffen Marburg and Bodo Nolte, editors, *Computational Acoustics of Noise Propagation in Fluids - Finite and Boundary Element Methods*, pp. 167–196. Springer Berlin Heidelberg, 2008.
- [BL85] C. K. Birdsall and A. B. Langdon. *Plasma Physics Via Computer Simulation*. The Adam Hilger series on plasma physics. McGraw-Hill, 1985.
- [BMG12] Didier Benisti, Olivier Morice, and Laurent Gremillet. “The various manifestations of collisionless dissipation in wave propagation.” *Physics of Plasmas*, **19**(6):063110, 2012.
- [Bor71] J. P. Boris. “Relativistic plasma simulation-optimization of a hybrid code.” In Jay P. Boris and Ramy A. Shanny, editors, *Proceedings of the Fourth Conference on Numerical Simulation of Plasmas*, pp. 3–67, Naval Research Laboratory, Washington, D.C., July 1971.
- [BPJ05] Lowell S. Brown, Dean L. Preston, and Robert L. Singleton Jr. “Charged particle motion in a highly ionized plasma.” *Physics Reports*, **410**(4):237 – 333, 2005.
- [BSG09] Didier Bénisti, David J. Strozzi, Laurent Gremillet, and Olivier Morice. “Nonlinear Landau Damping Rate of a Driven Plasma Wave.” *Physical Review Letters*, **103**(15):155002, Oct 2009.
- [BSS09] A. Bers, I. P. Shkarofsky, and M. Shoucri. “Relativistic Landau damping of electron plasma waves in stimulated Raman scattering.” *Physics of Plasmas*, **16**(2):022104, 2009.
- [BSW98] R. L. Berger, C. H. Still, E. A. Williams, and A. B. Langdon. “On the dominant and subdominant behavior of stimulated Raman and

- Brillouin scattering driven by nonuniform laser beams.” *Physics of Plasmas*, **5**(12):4337–4356, 1998.
- [Bun59] O. Buneman. “Dissipation of Currents in Ionized Media.” *Physical Review*, **115**(3):503–517, Aug 1959.
- [BV04] S. Brunner and E. J. Valeo. “Trapped-Particle Instability Leading to Bursting in Stimulated Raman Scattering Simulations.” *Physical Review Letters*, **93**(14):145003, Sep 2004.
- [BWL06] B. J. Albright, W. Daughton, L. Yin, K.J. Bowers, J.L. Kline, D.S. Montgomery, and J.C. Fernández. “Particle-in-cell studies of laser-driven hot spots and a statistical model for mesoscopic properties of Raman backscatter.” *Journal de Physique IV France*, **133**:253–257, 2006.
- [Che12] F. F. Chen. *Introduction to Plasma Physics*. Springer London, Limited, 2012.
- [Daw59] John M. Dawson. “Nonlinear Electron Oscillations in a Cold Plasma.” *Physical Review*, **113**(2):383–387, Jan 1959.
- [Daw83] John M. Dawson. “Particle simulation of plasmas.” *Reviews of Modern Physics*, **55**(2):403–447, Apr 1983.
- [Dec87] V. K. Decyk. “Simulation of microscopic processes in plasma.” In A. G. Sitenko, editor, *1987 International Conference on Plasma Physics. Joint Conference of the Seventh Kiev International Conference on Plasma Theory and the Seventh International Congress on Waves and Instabilities in Plasmas. Proceedings of the Invited Papers*, volume 2, pp. 1075–97, Singapore, Singapore, 1987. World Scientific.
- [DF99] C. Deutsch and P. Fromy. “Correlated stopping of relativistic electrons in superdense plasmas.” *Physics of Plasmas*, **6**(9):3597–3606, 1999.
- [DFL00] S. Depierreux, J. Fuchs, C. Labaune, A. Michard, H. A. Baldis, D. Pesme, S. Hüller, and G. Laval. “First Observation of Ion Acoustic Waves Produced by the Langmuir Decay Instability.” *Physical Review Letters*, **84**(13):2869–2872, Mar 2000.
- [DG67] D. F. DuBois and M. V. Goldman. “Parametrically Excited Plasma Fluctuations.” *Physical Review*, **164**(1):207–222, Dec 1967.
- [DH09] Jörn Dunkel and Peter Hänggi. “Relativistic Brownian motion.” *Physics Reports*, **471**(1):1 – 73, 2009.

- [DHH72] Ronald C. Davidson, David A. Hammer, Irving Haber, and Carl E. Wagner. “Nonlinear Development of Electromagnetic Instabilities in Anisotropic Plasmas.” *Physics of Fluids*, **15**(2):317–333, 1972.
- [DKL74] J. F. Drake, P. K. Kaw, Y. C. Lee, G. Schmid, C. S. Liu, and Marshall N. Rosenbluth. “Parametric instabilities of electromagnetic waves in plasmas.” *Physics of Fluids*, **17**(4):778–785, 1974.
- [DMD94] C. D. Decker, W. B. Mori, J. M. Dawson, and T. Katsouleas. “Nonlinear collisional absorption in laser-driven plasmas.” *Physics of Plasmas*, **1**(12):4043–4049, 1994.
- [DMO70] W. E. Drummond, J. H. Malmberg, T. M. O’Neil, and J. R. Thompson. “Nonlinear Development of the Beam-Plasma Instability.” *Physics of Fluids*, **13**(9):2422–2425, 1970.
- [DN04] Viktor K. Decyk and Charles D. Norton. “UCLA Parallel PIC Framework.” *Computer Physics Communications*, **164**(1-3):80 – 85, 2004. Proceedings of the 18th International Conference on the Numerical Simulation of Plasmas.
- [DS69] Heinrich Derfler and Thomas C. Simonen. “Higher-Order Landau Modes.” *Physics of Fluids*, **12**(2):269–278, 1969.
- [DTH10] E. L. Dewald, C. Thomas, S. Hunter, L. Divol, N. Meezan, S. H. Glenzer, L. J. Suter, E. Bond, J. L. Kline, J. Celeste, D. Bradley, P. Bell, R. L. Kauffman, J. Kilkeny, and O. L. Landen. “Hot electron measurements in ignition relevant Hohlraums on the National Ignition Facility.” *Review of Scientific Instruments*, **81**(10):10D938, 2010.
- [EK83] Kent Estabrook and W. L. Kruer. “Theory and simulation of one-dimensional Raman backward and forward scattering.” *Physics of Fluids*, **26**(7):1892–1903, 1983.
- [EKH89] Kent Estabrook, W. L. Kruer, and M. G. Haines. “Nonlinear features of stimulated Brillouin and Raman scattering.” *Physics of Fluids B: Plasma Physics*, **1**(6):1282–1286, 1989.
- [Fah10] J. E. Fahlen. *Nonlinear Phenomena of Plasma Waves in a Kinetic Regime: Frequency Shifts, Packets, and Transverse Localization*. PhD thesis, University of California, Los Angeles, 2010.
- [FC61] B. D. Fried and S. D. Conte. *The plasma dispersion function: the Hilbert transform of the Gaussian*. Academic Press, 1961.
- [FDL09] D. H. Froula, L. Divol, R. A. London, R. L. Berger, T. Döppner, N. B. Meezan, J. S. Ross, L. J. Suter, C. Sorce, and S. H. Glenzer. “Observation of the Density Threshold Behavior for the Onset of Stimulated

- Raman Scattering in High-Temperature Hohlraum Plasmas.” *Physical Review Letters*, **103**(4):045006, Jul 2009.
- [FKL75a] D. W. Forslund, J. M. Kindel, and E. L. Lindman. “Plasma simulation studies of stimulated scattering processes in laser-irradiated plasmas.” *Physics of Fluids*, **18**(8):1017–1030, 1975. And references therein.
- [FKL75b] D. W. Forslund, J. M. Kindel, and E. L. Lindman. “Theory of stimulated scattering processes in laser-irradiated plasmas.” *Physics of Fluids*, **18**(8):1002–1016, 1975.
- [FST02] 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 JC Adam. “OSIRIS: A Three-Dimensional, Fully Relativistic Particle-in-Cell Code for Modeling Plasma Based Accelerators.” In Peter Sloot, Alfons Hoekstra, C. Tan, and Jack Dongarra, editors, *Computational Science - ICCS 2002*, volume 2331 of *Lecture Notes in Computer Science*, pp. 342–351. Springer Berlin / Heidelberg, 2002. International Conference on Computational Science, Amsterdam, Netherlands, April 21-24, 2002.
- [FWG11] J. E. Fahlen, B. J. Winjum, T. Grismayer, and W. B. Mori. “Transverse plasma-wave localization in multiple dimensions.” *Physical Review E*, **83**(4):045401, Apr 2011.
- [GBB12] Frank R. Graziani, Victor S. Batista, Lorin X. Benedict, John I. Castor, Hui Chen, Sophia N. Chen, Chris A. Fichtl, James N. Glosli, Paul E. Grabowski, Alexander T. Graf, Stefan P. Hau-Riege, Andrew U. Hazi, Saad A. Khairallah, Liam Krauss, A. Bruce Langdon, Richard A. London, Andreas Markmann, Michael S. Murillo, David F. Richards, Howard A. Scott, Ronnie Shepherd, Liam G. Stanton, Fred H. Streitz, Michael P. Surh, Jon C. Weisheit, and Heather D. Whitley. “Large-scale molecular dynamics simulations of dense plasmas: The Cimarron Project.” *High Energy Density Physics*, **8**(1):105 – 131, 2012.
- [GCL04] Andrew D. Greenwood, Keith L. Cartwright, John W. Luginsland, and Ernest A. Baca. “On the elimination of numerical Cerenkov radiation in PIC simulations.” *Journal of Computational Physics*, **201**(2):665 – 684, 2004.
- [Gol84] Martin V. Goldman. “Strong turbulence of plasma waves.” *Reviews of Modern Physics*, **56**(4):709–735, Oct 1984.
- [GPS01] H. Goldstein, C. P. Poole, and J. L. Safko. *Classical Mechanics*. Addison-Wesley, New York, third edition, 2001.
- [GR71] K. W. Gentle and C. W. Roberson. “Observations of the Beam-Plasma Instability.” *Physics of Fluids*, **14**(12):2780–2782, 1971.

- [HDR06] C. Huang, V. K. Decyk, C. Ren, M. Zhou, W. Lu, W. B. Mori, J. H. Cooley, T. M. Antonsen Jr., and T. Katsouleas. “QUICKPIC: A highly efficient particle-in-cell code for modeling wakefield acceleration in plasmas.” *Journal of Computational Physics*, **217**(2):658 – 679, 2006.
- [HGE73] R. W. Hockney, S. P. Goel, and J. W. Eastwood. “A 10000 particle molecular dynamics model with long range forces.” *Chemical Physics Letters*, **21**(3):589 – 591, 1973.
- [HNU11] J. D. Huba, United States Office of Naval Research, and Naval Research Laboratory (U.S.). *NRL Plasma Formulary*. Naval Research Laboratory, 2011.
- [Jac99] J. D. Jackson. *Classical Electrodynamics*. Wiley, third edition, 1999.
- [JM07] Christopher S. Jones and Michael S. Murillo. “Analysis of semi-classical potentials for molecular dynamics and Monte Carlo simulations of warm dense matter.” *High Energy Density Physics*, **3**(34):379 – 394, 2007.
- [JNG05] Vince J. van Joolen, Beny Neta, and Dan Givoli. “High-order Higdon-like boundary conditions for exterior transient wave problems.” *International Journal for Numerical Methods in Engineering*, **63**(7):1041–1068, 2005.
- [KA75] Paul Koch and James Albritton. “Nonlinear Evolution of Stimulated Raman Backscatter in Cold Homogeneous Plasma.” *Physical Review Letters*, **34**(26):1616–1620, Jun 1975.
- [Kar07] M. Kardar. *Statistical Physics of Particles*. Cambridge University Press, 2007.
- [KD69] P. K. Kaw and J. M. Dawson. “Laser-Induced Anomalous Heating of a Plasma.” *Physics of Fluids*, **12**(12):2586–2591, 1969.
- [KEL80] W. L. Kruer, Kent Estabrook, B. F. Lasinski, and A. B. Langdon. “Raman backscatter in high temperature, inhomogeneous plasmas.” *Physics of Fluids*, **23**(7):1326–1329, 1980.
- [KM88] T. Katsouleas and W. B. Mori. “Wave-Breaking Amplitude of Relativistic Oscillations in a Thermal Plasma.” *Physical Review Letters*, **61**(1):90–93, Jul 1988.
- [KR71] Allan N. Kaufman and Peter S. Rostler. “The Darwin Model as a Tool for Electromagnetic Plasma Simulation.” *Physics of Fluids*, **14**(2):446–448, 1971.

- [Kru03] W. L. Kruer. *The Physics of Laser Plasma Interactions*. Frontiers in physics. Westview Press, 2003.
- [KT73] N. A. Krall and A. W. Trivelpiece. *Principles of plasma physics*. Number v. 0-911351 in International series in pure and applied physics. McGraw-Hill, 1973.
- [KWC87] T. Katsouleas, S. Wilks, P. Chen, J. M. Dawson, and J. J. Su. “Beam Loading in Plasma Accelerators.” *Particle Accelerators*, **22**:81–99, 1987.
- [LAB04] John D. Lindl, Peter Amendt, Richard L. Berger, S. Gail Glendinning, Siegfried H. Glenzer, Steven W. Haan, Robert L. Kauffman, Otto L. Landen, and Laurence J. Suter. “The physics basis for ignition using indirect-drive targets on the National Ignition Facility.” *Physics of Plasmas*, **11**(2):339–491, 2004.
- [Lan28] Irving Langmuir. “Oscillations in Ionized Gases.” *Proceedings of the National Academy of Sciences*, **14**(8):627–637, 1928.
- [Lan46] L. Landau. “On the vibrations of the electronic plasma.” *Journal of Physics (USSR)*, **10**(1):25–34, 1946. reproduced in D. ter Haar, ed., *Collected papers of L. D. Landau* (Pergamon Press, New York, 1965) p. 445-460, and in D. ter Haar, ed., *Men of Physics: L. D. Landau*, Vol. 2 (Pergamon Press, New York, 1965).
- [LCW07] R. R. Lindberg, A. E. Charman, and J. S. Wurtele. “Self-consistent Langmuir waves in resonantly driven thermal plasmas.” *Physics of Plasmas*, **14**(12):122103, 2007.
- [LL73] Roswell Lee and Martin Lampe. “Electromagnetic Instabilities, Filamentation, and Focusing of Relativistic Electron Beams.” *Physical Review Letters*, **31**(23):1390–1393, Dec 1973.
- [LLT13] R. Lehe, A. Lifschitz, C. Thauray, V. Malka, and X. Davoine. “Numerical growth of emittance in simulations of laser-wakefield acceleration.” *Physical Review Special Topics – Accelerators and Beams*, **16**(2):021301, Feb 2013.
- [LP06] C. K. Li and R. D. Petrasso. “Stopping, straggling, and blooming of directed energetic electrons in hydrogenic and arbitrary- Z plasmas.” *Physical Review E*, **73**(1):016402, Jan 2006.
- [LR76] C. S. Liu and Marshall N. Rosenbluth. “Parametric decay of electromagnetic waves into two plasmons and its consequences.” *Physics of Fluids*, **19**(7):967–971, 1976.

- [MA97] 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, 1997.
- [Mal99] S.G. Mallat. *A wavelet tour of signal processing*. Wavelet Analysis and Its Applications Series. Academic Press, 1999.
- [MAT12] *MATLAB Release 2012a*. The MathWorks Inc., Natick, Massachusetts, Feb 2012.
- [MBR09] E. I. Moses, R. N. Boyd, B. A. Remington, C. J. Keane, and R. Al-Ayat. “The National Ignition Facility: Ushering in a new age for high energy density science.” *Physics of Plasmas*, **16**(4):041006, 2009.
- [MCF02] D. S. Montgomery, J. A. Cobble, J. C. Fernandez, R. J. Focia, R. P. Johnson, N. Renard-LeGalloudec, H. A. Rose, and D. A. Russell. “Recent Trident single hot spot experiments: Evidence for kinetic effects, and observation of Langmuir decay instability cascade.” *Physics of Plasmas*, **9**(5):2311–2320, 2002.
- [MF02] V. M. Malkin and N. J. Fisch. “Collective Deceleration of Relativistic Electrons Precisely in the Core of an Inertial-Fusion Target.” *Physical Review Letters*, **89**(12):125004, Aug 2002.
- [MN98] Makoto Matsumoto and Takuji Nishimura. “Mersenne Twister: A 623-dimensionally Equidistributed Uniform Pseudo-random Number Generator.” *ACM Transactions on Modeling and Computer Simulation*, **8**(1):3–30, Jan 1998.
- [MO72] G. J. Morales and T. M. O’Neil. “Nonlinear Frequency Shift of an Electron Plasma Wave.” *Physical Review Letters*, **28**(7):417–420, Feb 1972.
- [Mor97] 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, Nov 1997.
- [MTE14] 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.
- [MWS07] Z. L. Mišković, Y. N. Wang, and Y. H. Song. “Dynamics of fast molecular ions in solids and plasmas: A review of recent theoretical developments.” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, **256**(1):57 – 65, 2007. Proceedings of the 22nd International Conference on Atomic Collisions in Solids.

- [NL76] Clair W. Nielson and H. Ralph Lewis. “Particle-Code Models in the Nonradiative Limit.” In John Killeen, editor, *Controlled Fusion*, volume 16 of *Methods in Computational Physics: Advances in Research and Applications*, pp. 367 – 388. Elsevier, 1976.
- [ON65] Thomas O’Neil. “Collisionless Damping of Nonlinear Plasma Oscillations.” *Physics of Fluids*, **8**(12):2255–2262, 1965.
- [OWM71] T. M. O’Neil, J. H. Winfrey, and J. H. Malmberg. “Nonlinear Interaction of a Small Cold Beam and a Plasma.” *Physics of Fluids*, **14**(6):1204–1212, 1971.
- [PRP10] J. P. Palastro, J. S. Ross, B. Pollock, L. Divol, D. H. Froula, and S. H. Glenzer. “Fully relativistic form factor for Thomson scattering.” *Physical Review E*, **81**(3):036411, Mar 2010.
- [PV91] Thomas Peter and Jürgen Meyer-ter Vehn. “Energy loss of heavy ions in dense plasma. I. Linear and nonlinear Vlasov theory for the stopping power.” *Physical Review A*, **43**(4):1998, 1991.
- [PW56] W. K. H. Panofsky and W. A. Wenzel. “Some Considerations Concerning the Transverse Deflection of Charged Particles in RadioFrequency Fields.” *Review of Scientific Instruments*, **27**(11):967–967, 1956.
- [Pyt12] *Python, Version 2.7.3*. Python Software Foundation, Apr 2012.
- [RGC09] D. F. Richards, J. N. Glosli, B. Chan, M. R. Dorr, E. W. Draeger, J. L. Fattebert, W. D. Krauss, T. Spelce, F. H. Streitz, M. P. Surh, and J. A. Gunnels. “Beyond homogeneous decomposition: scaling long-range forces on Massively Parallel Systems.” In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, SC ’09, pp. 60:1–60:12, New York, NY, USA, 2009. ACM.
- [RM84] International Commission on Radiation Units and Measurements. *Report 37, Stopping Powers for Electrons and Positrons*. 1984.
- [RR01] Harvey A. Rose and David A. Russell. “A self-consistent trapping model of driven electron plasma waves and limits on stimulated Raman scatter.” *Physics of Plasmas*, **8**(11):4784–4799, 2001.
- [RSC89] J. B. Rosenzweig, P. Schoessow, B. Cole, W. Gai, R. Konecny, J. Norem, and J. Simpson. “Experimental measurement of nonlinear plasma wake fields.” *Physical Review A*, **39**(3):1586–1589, Feb 1989.
- [RSC90] J. B. Rosenzweig, P. Schoessow, B. Cole, C. Ho, W. Gai, R. Konecny, S. Mtingwa, J. Norem, M. Rosing, and J. Simpson. “Demonstration of electron beam self-focusing in plasma wake fields.” *Physics of Fluids B: Plasma Physics*, **2**(6):1376–1383, 1990.

- [RTT06] C. Ren, M. Tzoufras, J. Tonge, W. B. Mori, F. S. Tsung, M. Fiore, R. A. Fonseca, L. O. Silva, J. C. Adam, and A. Heron. “A global simulation for laser-driven MeV electrons in 50- μ m-diameter fast ignition targets.” *Physics of Plasmas*, **13**(5):056308, 2006.
- [SB08] A. A. Solodov and R. Betti. “Stopping power and range of energetic electrons in dense plasmas of fast-ignition fusion targets.” *Physics of Plasmas*, **15**(4):042707, 2008.
- [SD05] K. V. Starikov and C. Deutsch. “Stopping of relativistic electrons in a partially degenerate electron fluid.” *Physical Review E*, **71**(2):026407, Feb 2005.
- [Sha63] V. D. Shapiro. “On Nonlinear Theory of Interaction Between a Monoenergetic Beam and a Plasma.” *Soviet Physics, Journal of Experimental and Theoretical Physics*, **17**(3):416, 1963.
- [SSA04] R. B. Stephens, R. A. Snavely, Y. Aglitskiy, F. Amiranoff, C. Andersen, D. Batani, S. D. Baton, T. Cowan, R. R. Freeman, T. Hall, S. P. Hatchett, J. M. Hill, M. H. Key, J. A. King, J. A. Koch, M. Koenig, A. J. MacKinnon, K. L. Lancaster, E. Martinolli, P. Norreys, E. Perelli-Cippo, M. Rabec Le Gloahec, C. Rousseaux, J. J. Santos, and F. Scianitti. “ K_α fluorescence measurement of relativistic electron transport in the context of fast ignition.” *Physical Review E*, **69**(6):066414, Jun 2004.
- [STL12] D. J. Strozzi, M. Tabak, D. J. Larson, L. Divol, A. J. Kemp, C. Bellei, M. M. Marinak, and M. H. Key. “Fast-ignition transport studies: Realistic electron source, integrated particle-in-cell and hydrodynamic modeling, imposed magnetic fields.” *Physics of Plasmas*, **19**(7):072711, 2012.
- [Str05] D. J. Strozzi. *Vlasov Simulations of Kinetic Enhancement of Raman Backscatter in Laser Fusion Plasmas*. PhD thesis, Massachusetts Institute of Technology, 2005.
- [SW72] R. Stenzel and A. Y. Wong. “Threshold and Saturation of the Parametric Decay Instability.” *Physical Review Letters*, **28**(5):274–277, Jan 1972.
- [SWH08] D. J. Strozzi, E. A. Williams, D. E. Hinkel, D. H. Froula, R. A. London, and D. A. Callahan. “Ray-based calculations of backscatter in laser fusion targets.” *Physics of Plasmas*, **15**(10):102703, 2008.
- [SWL07] D. J. Strozzi, E. A. Williams, A. B. Langdon, and A. Bers. “Kinetic enhancement of Raman backscatter, and electron acoustic Thomson scatter.” *Physics of Plasmas*, **14**(1):013104, 2007.

- [SWR12] D. J. Strozzi, E. A. Williams, H. A. Rose, D. E. Hinkel, A. B. Langdon, and J. W. Banks. “Threshold for electron trapping nonlinearity in Langmuir waves.” *Physics of Plasmas*, **19**(11):112306, 2012.
- [THG94] 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, 1994.
- [TL29] Lewi Tonks and Irving Langmuir. “Oscillations in Ionized Gases.” *Physical Review*, **33**(2):195–210, Feb 1929.
- [TM98] K. C. Tzeng and W. B. Mori. “Suppression of Electron Ponderomotive Blowout and Relativistic Self-Focusing by the Occurrence of Raman Scattering and Plasma Heating.” *Physical Review Letters*, **81**(1):104–107, Jul 1998.
- [TMM09] 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, 2009.
- [TRT06] M. Tzoufras, C. Ren, F. S. Tsung, J. W. Tonge, W. B. Mori, M. Fiore, R. A. Fonseca, and L. O. Silva. “Space-Charge Effects in the Current-Filamentation or Weibel Instability.” *Physical Review Letters*, **96**(10):105002, Mar 2006.
- [VDB01] H. X. Vu, D. F. DuBois, and B. Bezzerides. “Transient Enhancement and Detuning of Laser-Driven Parametric Instabilities by Particle Trapping.” *Physical Review Letters*, **86**(19):4306–4309, May 2001.
- [VDB02] H. X. Vu, D. F. DuBois, and B. Bezzerides. “Kinetic inflation of stimulated Raman backscatter in regimes of high linear Landau damping.” *Physics of Plasmas*, **9**(5):1745–1763, 2002.
- [VDB07] H. X. Vu, D. F. DuBois, and B. Bezzerides. “Inflation threshold: A nonlinear trapping-induced threshold for the rapid onset of stimulated Raman scattering from a single laser speckle.” *Physics of Plasmas*, **14**(1):012702, 2007.
- [Wei59] Erich S. Weibel. “Spontaneously Growing Transverse Waves in a Plasma Due to an Anisotropic Velocity Distribution.” *Physical Review Letters*, **2**(3):83–84, Feb 1959.
- [WFT10] B. J. Winjum, J. E. Fahlen, F. S. Tsung, and W. B. Mori. “Effects of plasma wave packets and local pump depletion in stimulated Raman scattering.” *Physical Review E*, **81**(4):045401, Apr 2010.

- [WFT13] B. J. Winjum, J. E. Fahlen, F. S. Tsung, and W. B. Mori. “Anomalous Hot Electrons due to Rescatter of Stimulated Raman Scattering in the Kinetic Regime.” *Physical Review Letters*, **110**(16):165001, Apr 2013.
- [Win10] B. J. Winjum. *Particle-In-Cell Simulations of Stimulated Raman Scattering for Parameters Relevant to Inertial Fusion Energy*. PhD thesis, University of California, Los Angeles, 2010.
- [Win12] Benjamin J. Winjum. Private communication, Jun 2012.
- [WKT92] 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.
- [WML09] T. L. Wang, D. Michta, R. R. Lindberg, A. E. Charman, S. F. Martins, and J. S. Wurtele. “Feasibility study for using an extended three-wave model to simulate plasma-based backward Raman amplification in one spatial dimension.” *Physics of Plasmas*, **16**(12):123110, 2009.
- [YAR12a] L. Yin, B. J. Albright, H. A. Rose, K. J. Bowers, B. Bergen, and R. K. Kirkwood. “Self-Organized Bursts of Coherent Stimulated Raman Scattering and Hot Electron Transport in Speckled Laser Plasma Media.” *Physical Review Letters*, **108**(24):245004, Jun 2012.
- [YAR12b] L. Yin, B. J. Albright, H. A. Rose, K. J. Bowers, B. Bergen, R. K. Kirkwood, D. E. Hinkel, A. B. Langdon, P. Michel, D. S. Montgomery, and J. L. Kline. “Trapping induced nonlinear behavior of backward stimulated Raman scattering in multi-speckled laser beams.” *Physics of Plasmas*, **19**(5):056304, 2012.
- [YF09] N. A. Yampolsky and N. J. Fisch. “Simplified model of nonlinear Landau damping.” *Physics of Plasmas*, **16**(7):072104, 2009.
- [YRL12] R. Yan, C. Ren, J. Li, A. V. Maximov, W. B. Mori, Z.-M. Sheng, and F. S. Tsung. “Generating energetic electrons through staged acceleration in the two-plasmon-decay instability in inertial confinement fusion.” *Physical Review Letters*, **108**(17):175002, Apr 2012.
- [Zak72] Vladimir E Zakharov. “Collapse of Langmuir waves.” *Soviet Physics, Journal of Experimental and Theoretical Physics*, **35**(5):908–914, 1972.