

Vlasov Simulation with FARSIGHT and Unlimited Photon Acceleration

by

Ryan Sandberg

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied and Interdisciplinary Mathematics and Scientific Computing)
in The University of Michigan
2022

Doctoral Committee:

Professor Robert Krasny, Co-Chair
Professor Alexander Thomas, Co-Chair
Assistant Professor Brendan Kochunas
Professor Shravan Veerapaneni

Ryan Sandberg

ryansand@umich.edu

ORCID iD: 0000-0001-7680-8733

© Ryan Sandberg 2022

All Rights Reserved

To Nicole, To Dad, To Mom – this degree belongs to all of you

ACKNOWLEDGEMENTS

I am grateful for my advisors, Robert Krasny and Alec Thomas, and the chance they took on me. They have given many hours of reading, listening, meeting, reading, guiding, and advising. As much as I've learned of numerical analysis and scientific investigation, I may have learned more from them of the human and artistic side of science and mathematics. We are storytellers and dreamers, adventurers and builders, and the people in science are what matter most.

I am grateful for my committee members, Brendan Kochunas and Shravan Veerapaneni. They have been in turn my teachers and my fans cheering me on through the final hurdles of my degree.

I am grateful for my wife, Nicole. She put aside her career and ambitions to be at home with our children and take care of me. She earned this degree as much as I did.

I am grateful for my parents. From before I can remember they have encouraged me to learn and supported me in my pursuit of knowledge.

I am grateful for all my family and for the many friends I've known in this time. There are too many to name and all have had a lasting impact.

This work was supported by AFOSR grant FA9550-19-1-0072. Computational resources were provided by Advanced Research Computing at the University of Michigan. Access to the OSIRIS 4.0 framework provided by the OSIRIS Consortium, consisting of UCLA and IST (Lisbon, Portugal). Work supported by NSF ACI-1339893.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vi
LIST OF TABLES	xii
ABSTRACT	xiii
 CHAPTER	
I. Introduction	1
1.1 Motivation	1
1.2 FARSIGHT	4
1.3 Unlimited photon acceleration	6
1.4 Summary	8
 II. Vlasov Simulation with FARSIGHT	 9
2.1 Introduction	9
2.1.1 Previous numerical methods	10
2.1.2 Present work	11
2.2 Preliminary topics	13
2.2.1 Flow map	13
2.2.2 Electric field integral	13
2.2.3 Kernel regularization	14
2.2.4 Electric field at a particle position	16
2.3 Numerical method	17
2.3.1 Phase-space discretization and adaptive mesh refinement	17
2.3.2 Storage of panel and particle information	18
2.3.3 Electric field quadrature and particle motion	19
2.3.4 Remeshing	21
2.3.5 Interpolation	23
2.3.6 Barycentric Lagrange treecode	24

2.3.7	GPU implementation	25
2.3.8	BLTC performance	26
2.4	Numerical Results	29
2.4.1	Landau damping	30
2.4.2	Two-stream instability	38
2.4.3	Cold two-stream instability	41
2.4.4	Halo formation in a mismatched thermal equilibrium sheet beam	42
2.5	Summary	47
III. Unlimited Photon Acceleration		50
3.1	Background : photon acceleration	52
3.2	Matching condition to mitigate dephasing	54
3.3	Description of the wake behind the driver	56
3.4	Description of the wake within the driver	63
3.4.1	Discussion of the wake solutions within the driver	71
3.5	Obtaining the phase matched density profile and frequency	72
3.6	Ultrashort and weak driver limits and unlimited photon acceleration	76
3.6.1	Ultrashort limit	76
3.6.2	Unlimited photon acceleration	77
3.6.3	Weak driver limit	80
3.7	3D considerations	81
3.8	Results	83
3.8.1	1D simulations with and without the PA^∞ density profile .	84
3.8.2	1D PA^∞ simulation for comparison with quasi-3D simulation	86
3.8.3	1D PA^∞ simulation with a long drive beam	86
3.8.4	1D PA^∞ with a piecewise approximation to the exact profile	89
3.8.5	Quasi-3D PA^∞ simulation with a short, broad drive beam .	91
3.8.6	Quasi-3D PA^∞ simulation with piecewise-density profile . .	94
3.9	Summary	96
3.9.1	Future work	97
IV. Summary		100
4.1	FARSIGHT	100
4.1.1	Future directions	101
4.2	Unlimited photon acceleration	103
4.2.1	Future directions	104
APPENDICES		108
BIBLIOGRAPHY		110

LIST OF FIGURES

Figure

1.1	Phase-space of the two-stream instability at times $t = 0, 30,$ and $60,$ showing the progressive filamentation and development of phase-space vortices that this dissertation seeks to resolve and understand. These images are generated from a FARSIGHT simulation.	4
1.2	FARSIGHT representation of phase-space as panels and their constituent 3×3 grid of particles.	5
1.3	Photon acceleration of an electromagnetic wave seeing a time-varying plasma density. In the special case of a density structure moving at the speed of light, the space-time variation becomes a gradient in the moving window coordinate $\zeta = t - z/c.$ In this figure there is an increasing gradient in index of refraction, indicated by the gradient in phase velocity $v_\phi.$ In the case of an underdense plasma, $v_\phi \sim 1 + 2n/n_c$ where $n_c = 4\pi^2 c^2 m_e \epsilon_0 / \lambda^2 e^2$ and being underdense means $n/n_c < 1.$ The dashed curve is the initial wave with wavelength λ_0 and the solid line shows the wave after some propagation time with wavelength $\lambda.$ The gradient in n and hence in v_ϕ results in the leading wave fronts moving less than the trailing wave fronts and a corresponding decrease in wavelength λ compared to $\lambda_0,$ i.e. $\lambda < \lambda_0.$	7
2.1	Elementary electric field kernel (a) sawtooth $k(s)$ in Eq. (2.8), splitting $k(s) = k_1(s) + k_2(s)$ in Eq. (2.9), (b) shifted sawtooth $k_1(s),$ (c) tophat $k_2(s),$ associated charge density (d) $k'(s),$ (e) $k'_1(s),$ (f) $k'_2(s),$ regularized elementary electric field kernel (g) $k_\epsilon(s)$ in Eq. (2.11), (h) shifted sawtooth $k_1(s),$ (i) regularized tophat $k_{2,\epsilon}(s),$ associated charge density (j) $k'_\epsilon(s),$ (k) $k'_1(s),$ (l) $k'_{2,\epsilon}(s).$	15
2.2	phase-space discretization by Lagrangian particles and panels, each panel is a rectangle with a 3×3 particle grid (\bullet), panels have a tree structure and are adaptively refined, (a) level 0, root panel comprising the entire computational domain, (b) level 1, uniformly refined in x and $v,$ (c) level 2, refined only in $v,$ (d) level 3, two panels refined in x and $v.$	18
2.3	Landau damping simulation by Eq. (2.16) illustrating how a fixed set of Lagrangian particles and panels is progressively deformed, (a) $t = 0,$ (b) $t = 3,$ (c) $t = 9.$	21

2.4	Remeshing, (a) regular particles/panels at initial time, (b) irregular particles/panels after evolving for some time, (c) new regular particles/panels (red), old irregular particles/panels (light blue), electron distribution is transferred from old particles to new particles using biquadratic interpolation in old panels, old particles/panels are then deleted and calculation proceeds with new particles/panels.	22
2.5	Treecode schematic, (a) the particles are partitioned into clusters (subintervals of $[0, L]$) with a hierarchical tree structure, in (b,c) the target particles (left, \bullet) interact with (b) source particles (right, \bullet) or with (c) Chebyshev points (right, \times) if there are enough source particles as determined by the MAC, cluster radius r , particle-cluster distance R	25
2.6	BLTC performance, single field evaluation, particles given by Eq. (2.23), error versus system size N , MAC $\theta = 0.6$, degree $n = 2 : 2 : 12$	27
2.7	BLTC performance, single field evaluation, particles given by Eq. (2.23), run time (s) versus error, $N = 1024^2$ particles, MAC $\theta = 0.2, 0.4, 0.6, 0.8$ constant on each curve, symbols give degree $n = 2, 4, 6, 8, 10, 12$ increasing from right to left on each curve, (a) 36 CPU cores, (b) one GPU, note different vertical scales in (a) and (b), horizontal lines give direct sum run time (s).	28
2.8	BLTC performance, single field evaluation, particles given by Eq. (2.23), MAC $\theta = 0.6$, degree $n = 8$ yielding error $\approx 10^{-6}$, run time (s) vs. system size N , direct sum and BLTC on 36 CPU cores and 1 GPU.	29
2.9	Weak Landau damping, 2-norm of electric field versus time, phase-space discretization $N = 256^2$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, dashed line gives predicted decay rate $\gamma = 0.1533$ [1], GPU run time 34.6 s.	30
2.10	Weak Landau damping, 2-norm of electric field versus time, effect of numerical parameters, top row: phase-space discretization $N = 32^2, 64^2, 128^2$, middle row: time step $\Delta t = 2, 1, 0.5$, bottom row: regularization parameter $\epsilon = 1.6, 0.8, 0.4$	31
2.11	Weak Landau damping, (a) initial electron distribution $f_0(x, v)$, (b)-(f) perturbation from Maxwellian, $\delta f(x, v, t) = f(x, v, t) - f_M(v)$, time $t \in [0, 60]$, phase-space discretization $N = 1024^2$, time step $\Delta t = 0.125$, regularization parameter $\epsilon = 0.05$, run time 1372 s.	32
2.12	Weak Landau damping, slices through electron density deviation from Maxwellian, $\delta f(0, v, t) = f(0, v, t) - f_M(v)$, time $t \in [0, 60]$, same parameters as Fig. 2.11.	33
2.13	Weak Landau damping, convergence of electron density with respect to numerical parameters, time $t = 60$, reference solution computed with phase-space discretization $N = 2048^2$, time step $\Delta t = 0.125$, regularization parameter $\epsilon = 0.05$, solid line shows asymptotic error scaling, (a) $N \in [32^2, 1024^2]$, error = $O(\Delta x^3)$, (b) $\Delta t \in [0.25, 2]$, error = $O(\Delta t^4)$, (c) $\epsilon \in [0.1, 0.8]$, error = $O(\epsilon^2)$	34

2.14	Weak Landau damping, variation in conserved quantities vs. time, effect of numerical parameters, (a) charge Q_N , (b) current I_N , (c) energy \mathcal{E}_N , phase-space discretization N , time step Δt , regularization parameter ϵ , blue (—): $(N, \Delta t, \epsilon) = (128^2, 1.0, 0.4)$, orange (—): $(N, \Delta t, \epsilon) = (256^2, 0.5, 0.2)$, green (—): $(N, \Delta t, \epsilon) = (512^2, 0.25, 0.1)$	34
2.15	Strong Landau damping, 2-norm of electric field vs. time, phase-space discretization $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, run time 754 s, decay rate ($\gamma = -0.2920$, orange line), growth rate ($\gamma = 0.0815$, black line) [2].	35
2.16	Strong Landau damping, electron distribution $f(x, v, t)$, discretization $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, run time 754 s.	36
2.17	Strong Landau damping, conservation properties, BLTC parameters MAC $\theta = 0.6$, degree $n = 8$, mesh size N , regularization parameter ϵ , time step Δt refined simultaneously, (a) total charge Q_N , (b) total current I_N , (c) total energy \mathcal{E}_N , (d) f_{\max} , (e) f_{\min} , (f) fraction of negative f , $ Q_{N-}/Q_N $, blue (—): $(N, \Delta t, \epsilon) = (128^2, 0.5, 0.8)$, orange (—): $(N, \Delta t, \epsilon) = (256^2, 0.25, 0.4)$, green (—): $(N, \Delta t, \epsilon) = (512^2, 0.125, 0.2)$, GPU times (s) 19.3, 156.9, 1356.4	37
2.18	Strong Landau damping, AMR simulation, $t_f = 60$, $N = 16 \times 64 + 5$ AMR levels, $N_0 = 154297$, $N_f = 538601$, or 14.7% to 51.2% of the number of points in the uniform calculation, comparable uniform resolution is $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, $\epsilon_{AMR} = 0.01$, GPU time 345 s, top: phase-space, bottom: panel level.	38
2.19	Strong Landau damping, AMR simulation, $N = 16 \times 64 + 5$ AMR levels, comparable uniform resolution is $N = 512 \times 2048$, regularization parameter $\epsilon = 0.1$, time step $\Delta t = 0.25$, AMR tolerance $\epsilon_{AMR} = 0.01$, run time 345 s, (a) phase-space at time $t = 60$, (b) zoom of small region in phase-space centered at $(x, v) = (2, 2)$, (c) phase-space panels in zoom region of (b). . .	39
2.20	Two-stream instability, phase-space electron distribution, initial condition Eq. (2.26), discretization $N = 16 \times 64 + 5$ AMR levels, $N = 264325$ initially, which is 25.1% of the uniform N and $N = 639165$ ultimately which is 60.8% of the uniform N , regularization parameter $\epsilon = 0.1$, run time 437 s compared to uniform run time 739 s.	40
2.21	Two-stream instability, time $t = 60$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, compare AMR with uniform resolution, (a) uniform discretization, $N = 528 \times 2048$, run time 739 s, (b) AMR, $\epsilon_{AMR} = 0.008$, $N = 16 \times 64 + 5$ AMR levels, equivalent to 25.1% to 60.8% of uniform N , run time 435 s, (c) AMR panels shaded by refinement level.	41
2.22	Cold two-stream instability, initial distribution Eq. (2.27), phase-space discretization $N = 16 \times 32 + 5$ AMR levels, $N_0 = 129245$ which is 24.6% of the uniform N , N_f is 193345, which is 36.8% of the uniform N , AMR tolerance $\epsilon_{AMR} = 0.008$, regularization parameter $\epsilon = 0.1$, time step $\Delta t = 0.25$, run time 44.1 s.	43

2.23	Cold two-stream instability, compare AMR with uniform resolution at time 22, $\epsilon = 0.1, \Delta t = 0.25$ (a) $N = 512 \times 1024$, GPU time 123.5 s, (b) $N = 16 \times 32 + 5$ levels of AMR, $\epsilon_{AMR} = 0.008$, GPU time 44.1 s, (c) AMR resolution.	44
2.24	Scaled density \tilde{n}_Δ for equilibrium sheet beam distribution [3, 4].	45
2.25	Mismatched beam problem, uniform phase-space discretization $N = 1024^2$, axial increment $\Delta s = 0.1L_p$, regularization parameter $\epsilon = 0.1$ mm, GPU run time 1169 s, results agree well with LTPIC reference solution in [4].	47
2.26	Mismatched beam problem, phase-space at axial position $s = 40L_p$, (a) uniform discretization, $N = 1024^2$, run time 1169 s, (b) AMR, $N = 64^2 + 6$ AMR levels, initially $N = 135253$ or 12.9% of N_u , finally $N = 207721$ or 19.8% of N_u , run time 245 s, (c) panels in AMR calculation.	48
3.1	Schematic of unlimited photon acceleration. (a) Line drawing of PA^∞ layout, drive beam (in orange) propagates to the right through the plasma, laser pulse (red) co-propagates at the point ζ_δ where density (theory in green, simulation in blue) is increasing with distance ζ from the leading edge of the drive beam. (b) 3D visualization of PA^∞ , drive beam in purple propagates to the right through the plasma, laser pulse in blue and red co-propagates, wake indicated with isocontours of plasma density depression (in green) and elevation (in yellow).	54
3.2	PA^∞ requires analytic expressions for position ζ_δ of zero perturbation and increasing plasma density and for the gradient there, $\left. \frac{\partial(\omega_p^2/\gamma)}{\partial\zeta} \right _{\zeta_\delta}$. These expressions are obtained by solving the plasma fluid equations behind the driver, the region indicated by the dashed black box here. The driver density is the orange line, creating the plasma response shown in blue. The analytic expression for the plasma response in the absence of the witness pulse is in green. Information about the fluid quantities at the end of the driver, where $\zeta = \zeta_d$, is also needed, but this is determined in section 3.4.	56
3.3	PA^∞ requires analytic expressions for plasma fluid quantities at the end of the driver, $\zeta = \zeta_d$. These expressions are obtained by solving the plasma fluid equations within the driver, the region indicated by the dashed black box in this figure. The driver density is the orange line, creating the plasma response within the driver shown in red (not to be confused with the oscillatory laser pulse further back in the wake, also in red).	64
3.4	Analytic expressions in the wake compared to OSIRIS 1D simulation data. Top row: $k_{p0}L_{d0} = 1.9, n_{d0}/n = 0.2, 0.5, 0.8$. Middle row: $k_{p0}L_{d0} = 6.28, n_{d0}/n = 0.1, 0.5, 0.6$ Bottom row: $k_{p0}L_{d0} = 9.5, n_{d0}/n = 0.1, 0.5, 0.6$. The expressions are exact in a uniform density plasma (here $n = 1$ but any value of n is possible).	73
3.5	Numerically determined profiles for unlimited photon acceleration, computed by algorithm 1 with $n_d = 0.4, L_d = 1$, and $A = 0.4$. (a) shows the computed plasma density and (b) shows the relative gain in laser frequency.	75

3.6	Comparing numerical and ultrashort small A expansions. We show that the wake amplitude γ_m , the position of 0 density perturbation ζ_δ , and the predicted density n and frequency gain k_L/k_{L0} profiles for $A = 0.01, 0.1, 1$, and 10 when $k_{p0}L_{d0} = 0.01$. The length satisfies $k_{p0}L_d \ll 1$ so $\gamma_m = 1 + A^2/2n$ is always accurate. The expansion to $\mathcal{O}(A^2)$ is quite accurate for ζ_δ, n , and k_L for A as large as 1 . The expansion when $A = 10$ is not good for ζ_δ but is surprisingly accurate in predicting n and k_L even for $A = 10$	82
3.7	Schematic of the PIC cycle.	85
3.8	Laser centroid in wake plotted over propagation distance, calculated with and without the PA^∞ tapered plasma density profile. Top panel: uniform density; the laser centroid dephases rapidly from the wake when the plasma density is uniform. Middle panel: PA^∞ tapered density profile; the laser centroid maintains its position relative to the wake. Bottom panel: piecewise-constant approximation to the tapered density profile; the laser maintains phase in the wake even in the case of only a piecewise approximation to the PA^∞ tapered density profile. To illustrate the wake we plot $\delta n = n_w(\xi; z) - n(z)$. Note that in these images the wake coordinate is $\xi = z - ct$ not ζ , but for $v_p \rightarrow c$, $\xi \simeq -c\zeta$	87
3.9	1D Results for a drive beam with density $n_d = 0.2$ and length $L_d = 1.9$ and a laser pulse with normalized strength parameter $a_0 = 2.4$, (a) spectrum of laser pulse versus propagation distance, the red curve shows the predicted frequency as determined by equations (3.76) and (3.77) and the black line shows the mean frequency, note that the mean frequency increases by a factor of 10, (b) Tapered density profile for dephasingless photon acceleration as determined by equations (3.76) and (3.77), (c) drive beam energy in blue and laser pulse energy in red, shown to different scales, the electron drive beam loses energy linearly to the wake and is almost completely depleted of energy, the laser gains about 7% of the initial energy in the drive beam . . .	88
3.10	1D results for a drive beam having density $n_d = 0.3$, length $L_d = 12.56$ and a laser pulse having normalized strength parameter $a_0 = 4.5$, (a) spectrum of laser pulse versus propagation distance, the red curve shows the predicted frequency as determined by equations (3.76) and (3.77) and the black line shows the mean frequency, note that the mean frequency increases by a factor of 30, (b) Tapered density profile for dephasingless photon acceleration as determined by equations (3.76) and (3.77), (c) drive beam energy in blue and laser pulse energy in red, shown to different scales, the electron drive beam loses energy linearly to the wake and loses about half its energy before parts of the beam are depleted and the wake is disrupted. Laser energy increases monotonically, with most of the energy gain occurring early in the simulation.	89

3.11	Results in 1D for drive beam $n_d = 0.3$, $L_d = 12.56$ and laser $a_0 = 4.5$, (a) spectrum of laser pulse versus propagation distance, measured in k_L/k_{L0} where $k_L = k_z$ is the laser wavenumber, the red curve shows the predicted frequency as determined by equations (3.76) and (3.77) and the black line shows the mean frequency, note that mean frequency increases by a factor of 30, (b) Tapered density profile (dashed red) for dephasingless photon acceleration as determined by equations (3.76) and (3.77) and piecewise approximation (solid blue), (c) drive beam energy in blue and laser pulse energy in red, shown to different scales, the electron drive beam loses energy linearly to the wake and loses about half its energy before parts of the beam are depleted and the wake is disrupted.	90
3.12	Isocontours from quasi-3d simulation data demonstrating phase matched photon acceleration; the purple disk at the right is the drive electron beam, yellow and green are the positive and negative density perturbations of the wake, and the blue and red spheroid is a contour of the envelope of the transverse electric field, colored to show the laser pulse field phase.	91
3.13	Density profile for unlimited photon acceleration driven by an ultrashort driver with $A = (k_{p0}L_{d0})(n_{d0}/n_0) = 0.38$	93
3.14	Quasi-3D simulation results from FBPIC for a drive beam with $k_{p0}L_{d0} = 0.5$. (a) Comparing the initial and final spectra. The spectrum has shifted to about 10× the initial wavenumber. (b) Plot of the spectrum over time. The red line is the theoretical 1D frequency shift. (c) Drive beam energy (blue) and witness laser energy (red)	94
3.15	Piecewise approximation to the profile used for quasi-3D simulation	95
3.16	Quasi-3D simulation results from FBPIC for a drive beam with $k_{p0}L_{d0} = 0.5$ using a piecewise approximation to the phase matched density profile. (a) Comparing the initial and late-time ($z = 72$ mm) spectra. The spectrum has shifted to almost 10x the initial wavenumber. (b) Plot of the spectrum over time. The red line is the theoretical 1D frequency shift. (c) Drive beam energy (blue) and witness laser energy (red)	96
4.1	(a) Relative frequency gain as a function of density, $(\omega/\omega_0)(n)$, is plotted for beam drive charges $Q = 1$ nC, 2 nC, and 5 nC. The frequency gain-density curves are labeled with markers indicating the distance that the laser and drive beam must propagate to achieve the indicated shift. (b) Relative frequency gain as a function of z in meters, plotted for beam drive charges $Q = 1$ nC, 2 nC, and 5 nC.	106

LIST OF TABLES

Table

2.1	Panel data, (a) data associated with each panel, (b) representation of data.	19
2.2	Parameters for equilibrium sheet beam distribution from Campos Pinto et al. [4].	45
2.3	Mismatched beam profile, profile of results in Fig. 2.26, discretization N , initial and final values N_0 , N_f , average value N_{ave} , total run time, field evaluation time, remeshing time	48

ABSTRACT

This dissertation presents two bodies of work on the theoretical and numerical study of plasma waves via kinetic models, resolving the fine-scale structures that arise in phase space, and the implications for plasma-based particle acceleration.

First, we present the plasma simulation method FARSIGHT, a first step toward solving problems in relativistic plasma physics such as photon acceleration. It is a forward semi-Lagrangian particle method for the Vlasov-Poisson system in which the particle number density is represented on adaptively refined and remeshed panels in phase space, and an integral form of the Poisson equation is solved using a regularized electric field kernel and a GPU-accelerated hierarchical treecode. We describe the method and implementation and present numerical results encompassing Landau damping, two-stream instability, and halo formation in a particle beam. These results show the method's ability to resolve fine-scale features in phase space.

Second, we present unlimited photon acceleration, (PA^∞), a scheme for dephasingless photon acceleration in a particle-beam-driven wake. Photon acceleration and deceleration occurs when electromagnetic radiation experiences a time-varying plasma density; the time-varying density causes frequency changes in the laser pulse. In PA^∞ , a relativistic electron bunch propagates through a plasma and leaves density variations in its wake. A laser pulse propagates behind the electron beam so as to experience a region of frequency-upshifting density variation in the wake. Using a tapered density profile to keep the laser pulse at the phase in the wake where the frequency of the pulse is increasing, simulations suggest that the

laser pulse can see sustained frequency shift, as well as energy gain, intensity enhancement, and pulse compression. In one dimension, the frequency increases $25\times$, energy $6\times$, intensity $25\times$, and compression $33\times$. In quasi-3d simulations, the frequency increases $10\times$, energy $5\times$, intensity $20\times$, and compression $3\times$. It is mathematically demonstrated that the frequency shift of the laser pulse is limited only by the ability to maintain the wake, that is, the photon acceleration is unlimited.

CHAPTER I

Introduction

This dissertation presents two bodies of work on the theoretical and numerical study of plasma waves via kinetic models, resolving the fine-scale structures that arise in phase-space, and the implications for plasma-based particle acceleration. Chapter I introduces the work. Chapter II presents an adaptive semi-Lagrangian method for the Vlasov-Poisson system inspired by vortex methods in fluid dynamics; the method is called FARSIGHT. Chapter III presents a scheme for frequency shifting an electromagnetic wave in a tailored plasma density profile. Frequency upshift with this tailored profile is one of the applications eventually intended for the semi-Lagrangian method FARSIGHT. In this chapter we present motivation for the work, followed by descriptions of the remaining chapters.

1.1 Motivation

Particle beams and intense radiation have a myriad of uses, including imaging, beam therapy, photolithography for microchip manufacture, space thrusters, high-energy-physics, material science, generation of warm, dense matter, heating in nuclear fusion applications, and as objects of study in their own right. Fundamental particle physics uses kilometer-scale accelerators and high-energy particle beams to probe the fundamental structure of matter. Accelerators and radiation sources across a range of scales are used in imaging

technology for the study of biological and chemical systems and material science and in medicine as radiography, imaging, and beam therapy. Conventional accelerators are limited to accelerating gradients of about 100 MV/rmm [5], beyond which the material surface of the accelerator ionizes and breaks down. On the other hand, a plasma is already broken down and can sustain accelerating gradients of 100 GV/m or more, 1000× greater than the limiting gradients in conventional accelerators. Thus plasma-based particle acceleration schemes can accomplish comparable acceleration in a fraction of the size and cost of conventional accelerators. Plasma-based schemes are actively being investigated as candidates for the next generation of particle colliders [6] and as sources of radiation [7]. Analogous to plasma-based particle acceleration, the same plasma structures that can accelerate particles to high energies can cause frequency shifts and energy gains of electromagnetic waves propagating in plasmas. To highlight the analogy, the phenomenon was termed ‘photon acceleration’ [8].

Particle kinetics are essential to understanding plasma-based acceleration of electrons and photons. A kinetic model describes the evolution of the particle distribution function $f_s(\mathbf{x}, \mathbf{p}, t)$ for a species s where f_s is the probability distribution of species s taking values at position \mathbf{x} and momentum \mathbf{p} at time t . For massive particles, the momentum is $\mathbf{p} = \gamma m_s \mathbf{v}$ where m_s is the species mass, \mathbf{v} is the particle velocity, and $\gamma = (1 - \mathbf{v}^2/c^2)^{1/2}$ is the Lorentz factor. Analogously, a frequency-time representation describes the position-frequency evolution of an electromagnetic wave. In the short wavelength, slowly changing envelope limit of geometric optics we can treat the light as photons with frequency ω and wavenumber \mathbf{k} , in which case $\mathbf{v} = \omega/k$ is the phase velocity of a photon in the medium. In each case the details of the accelerating structure can be understood with a kinetic model described by the Vlasov equation [9],

$$\frac{\partial f_s(\mathbf{x}, \mathbf{p}, t)}{\partial t} + \sum_{i=1}^3 v_i \frac{\partial f_s(\mathbf{x}, \mathbf{p}, t)}{\partial x_i} + \sum_{i=1}^3 F_i(\mathbf{x}, \mathbf{p}, t) \frac{\partial f_s(\mathbf{x}, \mathbf{p}, t)}{\partial p_i} = 0 . \quad (1.1)$$

For charged particle species s with charge q_s , the force \mathbf{F} is given by the Lorentz force $q_s(\mathbf{E} + \mathbf{v} \times \mathbf{B})$ and the electromagnetic fields \mathbf{E}, \mathbf{B} are determined by Maxwell's equations. For photons with frequency $\omega(\mathbf{x}, \mathbf{p}, t)$ in the limit of geometric optics, $\mathbf{F} = -\nabla\omega$ [10].

There are several numerical challenges to solving the Vlasov equation. Even for the electrostatic Vlasov-Poisson system, these include:

- the general Vlasov equations are 7 dimensional, having one time, three position, and three momentum coordinates,
- physical quantities that should be conserved, such as mass/charge, momentum, and energy, are not dependent variables but are moments of the distribution function,
- the distribution function admits evolution at a large range of momenta and so multiple time scales must be resolved,
- the collisionless nature of the Vlasov equation means that phase-space structures can develop on increasingly smaller spatial scales,
- multiple plasma species with disparate charge-to-mass ratios can introduce even more time scales that have to be resolved.

The phase-space structures mentioned in the penultimate point include phase-space vortices arising in the simplest plasma systems, for example in periodic 1D Vlasov-Poisson, and throughout many branches of plasma physics, including inertial confinement fusion (cross-beam energy transfer [11]), astrophysics (shock acceleration [12, 13]), and space physics (electron holes in the solar wind [14]), to name but a few examples. Consider the phase-space shown in figure 1.1. In addition to horizontal layering, there is an intricate mixing of high- and low-density distribution regions in phase-space vortices. This dissertation seeks to accurately model and use these complex phase-space structures for applications such as particle acceleration.




Figure 1.1: Phase-space of the two-stream instability at times $t = 0$, 30, and 60, showing the progressive filamentation and development of phase-space vortices that this dissertation seeks to resolve and understand. These images are generated from a FARSIGHT simulation.

1.2 FARSIGHT

Chapter II of this dissertation presents a numerical method called FARSIGHT for solving the Vlasov-Poisson system. The motivation for FARSIGHT is to develop efficient adaptive refinement of phase-space to resolve fine-scale structures of kinetic plasma models.

FARSIGHT uses a kinetic model of a single-component plasma described by the normalized Vlasov-Poisson system

$$\partial_t f + v \partial_x f - E \partial_v f = 0, \quad E = -\partial_x \phi, \quad -\partial_x^2 \phi = \rho, \quad (1.2)$$

where $f(x, v, t)$ is the particle distribution function in (x, v) phase-space, with electric field $E(x, t)$, potential $\phi(x, t)$, and charge density $\rho(x, t)$. This is the 1D nonrelativistic electrostatic version of eqn. (1.1) where velocity, v , is used instead of momentum, p . It is normalized by the particle charge-to-mass ratio. Periodic boundary conditions are enforced on $0 \leq x \leq L$, and the charge density is

$$\rho(x, t) = 1 - \int_{-\infty}^{\infty} f(x, v, t) dv, \quad \int_{-\infty}^{\infty} f(x, v, t) dx = 1. \quad (1.3)$$

One physical interpretation of this single-component plasma model is a mobile electron species in a static uniform distribution of positive ions.

A semi-Lagrangian approach combines the accuracy advantages of Lagrangian particle tracking with the smoothness of a grid-based Eulerian representation of the particle distribution. A hierarchical tree of quadrilateral panels is used to represent the distribution, as shown in figure 1.2. Each panel consists of a 3×3 grid of Lagrangian particles, with particle charge determined from the panel area and the particle's distribution value. The panel structure, including some adaptively refined panels, is shown in Fig. 1.2. The particles are advected forward in time according to Newton's equations of motion. To solve Poisson's equation for the electric potential, the charge density is convolved with a regularized Green's function and the potential is then differentiated to obtain the electric field, which can be understood as a convolution of a regularized electric field kernel with the charge density. The convolution integral is computed by summing over the tree of panels. The quadrature scheme loses accuracy as the mesh distorts, so remeshing is applied. The distribution is interpolated from its values on the distorted mesh to a regular mesh using biquadratic interpolation. When remeshing, panels are adaptively refined whenever the variation in distribution is too large in a panel. The particles then continue to advect.




Figure 1.2: FARSIGHT representation of phase-space as panels and their constituent 3×3 grid of particles.

Several features of FARSIGHT enable efficient adaptive refinement of phase-space to resolve fine-scale phase-space structures. The panels in the hierarchical tree are adaptively refined with a local refinement scheme. This promotes parallel efficiency as the leaves are independent for purpose of field evaluation and interpolation. Another feature promoting efficient resolution of phase-space is the regularization of the electric field kernel, which allows larger time steps and, hence, fewer field evaluations. A third feature is the use of a treecode. In a system with N particles, the integral field evaluation requires $\mathcal{O}(N^2)$ operations via direct summation. In this work the cost is reduced to $\mathcal{O}(N \log N)$ with a GPU-accelerated treecode for further computational speedup.

In the implementation of FARSIGHT, the GPU was found to give a speedup of $10\times$ over a 36 core CPU compute node. The FARSIGHT code is then demonstrated on several test cases, including Landau damping, the two-stream instability, and halo formation in a mismatched thermal equilibrium sheet beam in a continuous focusing channel. These examples highlight the scheme's ability to maintain resolution over long time. In these examples, the use of adaptive refinement leads to as much as $4\times$ speedup over uniform mesh calculations.

1.3 Unlimited photon acceleration

The second work on plasma waves and plasma-based acceleration treated in this dissertation is unlimited photon acceleration. The concept of photon acceleration involves frequency shift of an electromagnetic wave propagating in a plasma with a time-varying plasma density. The shift can be understood as follows. In moving window coordinates $z, \zeta = t - z/c$, time variation becomes variation in ζ . Consider an underdense plasma, meaning that the plasma density n is much less than the critical density n_c of the electromagnetic wave. The critical density is given by $n_c = 4\pi^2 c^2 m_e \epsilon_0 / \lambda^2 e^2$ where λ is the wavelength of the electromagnetic wave. If the electromagnetic wave experiences a positive gradient in phase velocity with

respect to ζ , as illustrated in figure 1.3, then leading wavefronts move slower than trailing wavefronts and the wavelength λ decreases. For frequency $\omega^2 = \omega_p^2 + (2\pi c/\lambda)^2$ where the plasma frequency is defined by $\omega_p^2 = ne/m_e\epsilon_0$ and the phase velocity is v_ϕ , photon acceleration can be expressed by the frequency shift relation [15, 16],

$$\frac{1}{\omega} \frac{\partial \omega}{\partial z} = \frac{\partial v_\phi}{\partial \zeta}. \quad (1.4)$$

The phase velocity of the wave is increasing in ζ , which means leading wave fronts have a slower phase velocity so the trailing wavefronts catch up. Then the wavelength is compressed and the frequency increases.




Figure 1.3: Photon acceleration of an electromagnetic wave seeing a time-varying plasma density. In the special case of a density structure moving at the speed of light, the space-time variation becomes a gradient in the moving window coordinate $\zeta = t - z/c$. In this figure there is an increasing gradient in index of refraction, indicated by the gradient in phase velocity v_ϕ . In the case of an underdense plasma, $v_\phi \sim 1 + 2n/n_c$ where $n_c = 4\pi^2 c^2 m_e \epsilon_0 / \lambda^2 e^2$ and being underdense means $n/n_c < 1$. The dashed curve is the initial wave with wavelength λ_0 and the solid line shows the wave after some propagation time with wavelength λ . The gradient in n and hence in v_ϕ results in the leading wave fronts moving less than the trailing wave fronts and a corresponding decrease in wavelength λ compared to λ_0 , i.e. $\lambda < \lambda_0$.

Photon acceleration was first presented in 1989 by Wilks et al. [8], wherein 10% frequency shifts were observed in simulation. The shift was analyzed more rigorously by Esarey et al. [17] and proposed theoretical limits on possible frequency shifts presented by Mironov et

al. [18]. Photon acceleration/deceleration has been explored as a wakefield diagnostic [19]. Related phenomena for generating frequency shifts have been recently studied and frequency shifts of as much as $10\times$ have been considered, but energy retention frequency was found to decrease with increasing frequency shift [20, 21, 22].

Photon acceleration has not yet been realized as a mechanism for generating significant frequency shift. One substantial issue to overcome is the dephasing of the electromagnetic wave from the point in the wake where positive frequency shift occurs. To overcome dephasing, we introduce a tailored plasma density profile to match the wake phase of the co-propagating laser pulse with the frequency shifting region of the wake. Simulations show that more than $10\times$ frequency shift can be realized by a laser pulse propagating behind an ultrarelativistic electron beam through the tailored plasma density profile. In chapter III we present the frequency shift relations, derive the phase matching conditions, and present numerical studies of phase matched plasma wakefield photon acceleration in 1D and quasi-3D, realizing $10\times$ frequency shift, $5\times$ energy gain, $20\times$ increase in intensity of the laser pulse, and $3\times$ increase in laser compression.

1.4 Summary

Chapter IV summarizes the FARSIGHT method and results and unlimited photon acceleration. Future directions for each topic are discussed.

CHAPTER II

Vlasov Simulation with FARSIGHT

This chapter presents the FARSIGHT method for Vlasov simulation. The rest of this section, section 2.1, introduces the Vlasov-Poisson system of equations, discusses previous numerical methods for solving the Vlasov-Poisson system, and outlines some key features of FARSIGHT and how FARSIGHT is distinguished from other methods. Section 2.2 presents several preliminary topics including the flow map, electric field integral, kernel regularization, and the electric field following a particle. Section 2.3 presents the details of the FARSIGHT numerical method. Section 2.4 presents numerical results for weak and strong Landau damping, warm and cold two-stream instability, and halo formation for a mismatched thermal equilibrium sheet beam in a focusing channel [4]. Section 2.5 summarizes the chapter and outlines directions of future work in FARSIGHT.

2.1 Introduction

Collisionless plasmas modeled using the 1D1V Vlasov-Poisson system of equations arise in many applications including shock acceleration [12, 13], beam physics [3, 4], inertial confinement fusion [23, 24], and plasma-wakefield acceleration [25, 26]. Here we consider the kinetic model of a one-component electron plasma described by the 1D1V Vlasov-Poisson

equations in non-dimensional form,

$$\partial_t f + v \partial_x f - E \partial_v f = 0, \quad E = -\partial_x \phi, \quad -\partial_x^2 \phi = \rho, \quad (2.1)$$

where $f(x, v, t)$ is the electron distribution function in (x, v) phase-space, with electric field $E(x, t)$, potential $\phi(x, t)$, and charge density $\rho(x, t)$. Periodic boundary conditions are enforced on $0 \leq x \leq L$, and the charge density is

$$\rho(x, t) = 1 - \int_{-\infty}^{\infty} f(x, v, t) dv, \quad \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, v, t) dx dv = 1, \quad (2.2)$$

where a static, uniform distribution of charge is present to satisfy neutrality. We will utilize this formulation to simulate Landau damping and two-stream instability; we will also simulate halo formation for a mismatched thermal equilibrium sheet beam in a focusing channel [4], and the changes to the formulation needed in that case will be given later on. To set the context for the present work, next we give a brief partial overview of numerical methods for plasma simulations.

2.1.1 Previous numerical methods

The conventional choice for kinetic plasma simulations is the particle-in-cell (PIC) method in which the plasma is represented by a finite set of particles in phase-space and the electric field is computed using a grid-based Maxwell or Poisson solver [27, 28]. PIC is in widespread use, but difficulties may arise due to discrete particle noise unless sufficiently many particles are used [4, 29]. In addition, particle-grid interactions may induce thermal heating and energy drift [27, 28]. Several methods were developed to address these issues within the PIC framework including Complex Particle Kinetics (CPK) using Gaussian-shaped particles that can merge or fragment [30], Linearly Transformed PIC (LTPIC) using deformation matrices to give each particle a finite width in velocity [4], and Shape Function Kinetics (SFK) using several particles to track the local deformation of the phase-space distribution [31]. For elec-

trostatic PIC simulations with periodic boundary conditions, the potential and electric field are computed using fast grid-based Poisson solvers such as the FFT, but boundary integral methods have also been employed for problems with complex geometry [32, 33, 34].

Alternatively, Eulerian methods were developed that directly solve the Vlasov equation on a grid or using a spectral representation of the distribution function. These representations are typically smoother than those provided by PIC, and this is beneficial for example in simulating halo formation in particle beams [23, 4] and the dynamics of electron holes [35, 36]. Eulerian approaches include finite volume methods [37, 38] finite element and discontinuous Galerkin methods [39], and spectral methods [40, 41, 26].

Eulerian methods, however, may require small time steps for stability, and this is alleviated in semi-Lagrangian methods that permit large time steps by combining particle tracking with a grid-based representation of the distribution function [42, 43, 44]. In these codes the spatial discretization is often done using discontinuous Galerkin methods [1, 45, 46] or weighted essentially non-oscillatory (WENO) methods [2, 47, 48]. While early semi-Lagrangian methods used operator splitting [43], multidimensional methods have been employed recently [49, 50]. Since uniform grids are inefficient for problems with low particle density regions in phase-space, another recent trend is the use of spatially adaptive methods; the techniques employed include adaptive mesh refinement [51, 52], multi-grid remeshing [53], block decomposition [54] and multiresolution wavelet representations [55, 56, 57].

2.1.2 Present work

We present a forward semi-Lagrangian scheme called FARSIGHT for collisionless electrostatic plasmas described by the 1D1V Vlasov-Poisson equations. The phase-space distribution function is represented by quadrilateral panels having a hierarchical tree structure, and each panel is a 3×3 grid of particles that is evolved by 4th order Runge-Kutta time-stepping. The electric field is expressed as a convolution integral of the charge density with a regu-

larized electric field kernel. The particles are remeshed at every time step using biquadratic interpolation on the panels, and the panels are adaptively refined to resolve the phase-space distribution function. The electric field integral is discretized by the trapezoid rule, and the discrete sums are computed by a GPU-accelerated barycentric Lagrange treecode [58].

A few comments follow to explain this choice of techniques: the overriding goal was efficient, adaptive, high-fidelity modeling of phase-space. Previous work in fluid mechanics advantageously represented flow on the surface of a sphere with adaptively refined triangular or quadrilateral panels [59]; this success motivated the choice in developing FARSIGHT to represent the distribution function by a hierarchical tree of adaptively refined quadrilateral panels. The electric field is expressed as an integral and computed by summing over the leaf panels in the tree; this promotes parallel efficiency because the leaves are independent of each other, and it avoids the artificial diffusion and communication costs that may arise in discretizing and solving the Poisson equation on a regular grid. The electric field kernel has a discontinuity which obstructs accurate particle tracking, so the kernel is regularized to overcome this problem. In a system with N particles, the integral field solver requires $\mathcal{O}(N^2)$ operations using direct summation, but the cost is reduced to $\mathcal{O}(N \log N)$ using the treecode. The barycentric Lagrange version of the treecode used here is kernel-independent, so it can handle the regularized electric field kernel, and it is GPU-accelerated for efficiency.

The outline of this chapter is as follows. Section 2.2 discusses several preliminary topics including the flow map, electric field integral, kernel regularization, and the electric field following a particle. Section 2.3 develops the numerical method. Section 2.4 presents numerical results for weak and strong Landau damping, warm and cold two-stream instability, and halo formation for a mismatched thermal equilibrium sheet beam in a focusing channel [4]. Section 2.5 summarizes the work.

2.2 Preliminary topics

This section discusses several preliminary topics needed for the numerical method including the flow map, electric field integral, kernel regularization, and the electric field at a particle position.

2.2.1 Flow map

The particle flow map, $x(x_0, v_0, t), v(x_0, v_0, t)$, gives the location of a particle in phase-space at time t , labeled by its initial coordinates x_0, v_0 . The governing equations are

$$\partial_t x(x_0, v_0, t) = v(x_0, v_0, t), \quad \partial_t v(x_0, v_0, t) = -E(x(x_0, v_0, t), t), \quad (2.3)$$

where $E(x, t)$ is the electric field. It follows from the Vlasov equation that the distribution function is invariant in the phase flow,

$$f(x(x_0, v_0, t), v(x_0, v_0, t), t) = f_0(x_0, v_0), \quad (2.4)$$

where $f_0(x_0, v_0)$ is the initial distribution. The Vlasov-Poisson equations conserve several quantities including the total charge Q , total current I , and total energy \mathcal{E} ,

$$Q = \int_0^L \int_{-\infty}^{\infty} f dv dx, \quad I = \int_0^L \int_{-\infty}^{\infty} v f dv dx, \quad \mathcal{E} = \frac{1}{2} \int_0^L \int_{-\infty}^{\infty} v^2 f dv dx + \frac{1}{2} \int_0^L E^2 dx. \quad (2.5)$$

2.2.2 Electric field integral

The Poisson equation for the electric potential can be solved by convolving the charge density with the 1D periodic Laplace Green's function G ,

$$\phi(x, t) = - \int_0^L \int_{-\infty}^{\infty} G(x, y) f(y, v, t) dv dy, \quad (2.6)$$

and then negating and differentiating the resulting expression, which yields the electric field integral,

$$E(x, t) = \int_0^L \int_{-\infty}^{\infty} K(x, y) f(y, v, t) dv dy, \quad (2.7)$$

where the electric field kernel is

$$K(x, y) = k\left(\frac{x - y}{L}\right), \quad k(s) = \text{mod}(s, 1) - \frac{1}{2}, \quad (2.8)$$

and the elementary kernel $k(s)$ has the sawtooth form in Fig. 2.1a. Note that $k(s)$ is 1-periodic ($k(s+1) = k(s)$), piecewise linear with slope 1 between the integers $s = 0, \pm 1, \pm 2, \dots$, and has a negative jump discontinuity at the integers; hence the associated charge density $k'(s)$ is a periodic array of negative delta functions with a uniform neutralizing positive background (Fig. 2.1d).

2.2.3 Kernel regularization

The discontinuity in $k(s)$ implies a discontinuity in the electric field kernel $K(x, y)$ at $x = y$, which is problematic for discretizing the electric field integral and numerically solving the flow map equations in Eq. (2.3); in particular, when particles cross each other in physical space, the discontinuous force makes it difficult to obtain high accuracy unless excessively small time steps are used. To overcome this we follow the approach taken in vortex methods for incompressible fluid flow in which the kernel is regularized [60, 61, 62]. The regularization presented here is novel, to the best of our knowledge. We start by writing the elementary kernel as the sum of two terms,

$$k(s) = k_1(s) + k_2(s), \quad k_1(s) = k\left(s - \frac{1}{2}\right), \quad k_2(s) = -\frac{1}{2}\text{sign}(k_1(s)), \quad (2.9)$$

i.e., $k_1(s)$ is the shifted sawtooth shown in Fig. 2.1b and $k_2(s)$ is the tophat function shown in Fig. 2.1c. Note that $k_1(s)$ can be understood as the field due to a periodic array of

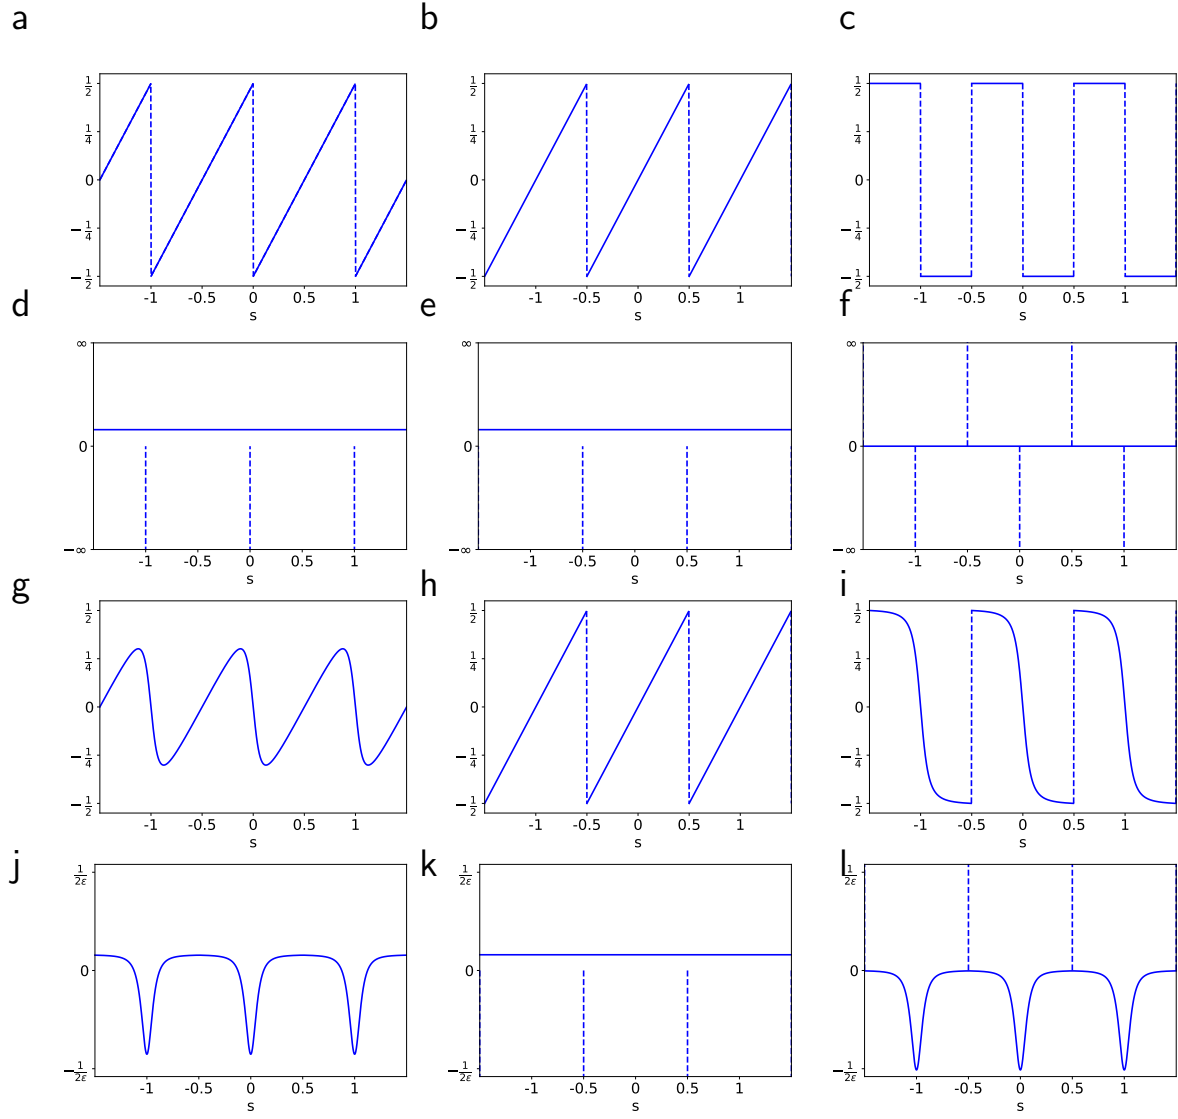


Figure 2.1: Elementary electric field kernel (a) sawtooth $k(s)$ in Eq. (2.8), splitting $k(s) = k_1(s) + k_2(s)$ in Eq. (2.9), (b) shifted sawtooth $k_1(s)$, (c) tophat $k_2(s)$, associated charge density (d) $k'(s)$, (e) $k_1'(s)$, (f) $k_2'(s)$, regularized elementary electric field kernel (g) $k_\epsilon(s)$ in Eq. (2.11), (h) shifted sawtooth $k_1(s)$, (i) regularized tophat $k_{2,\epsilon}(s)$, associated charge density (j) $k'_\epsilon(s)$, (k) $k_1'(s)$, (l) $k_{2,\epsilon}'(s)$.

negative discrete charges at the half-integers with a uniform neutralizing positive background (Fig. 2.1e), while $k_2(s)$ can be understood as the field due to two periodic arrays, negative discrete charges at the integers and positive discrete charges at the half-integers (Fig. 2.1f).

Based on this splitting we define the regularized elementary kernel,

$$k_\epsilon(s) = k_1(s) + k_{2,\epsilon}(s), \quad k_{2,\epsilon}(s) = -\frac{c_\epsilon}{2} \frac{k_1(s)}{(k_1(s)^2 + \epsilon^2/L^2)^{1/2}}, \quad c_\epsilon = (1 + 4\epsilon^2/L^2)^{1/2}, \quad (2.10)$$

where $\epsilon > 0$ is the regularization parameter. Figure 2.1g depicts $k_\epsilon(s)$ for a representative value of ϵ , and Figs. 2.1h,i show the two components, where $k_1(s)$ has not changed and $k_{2,\epsilon}(s)$ is a regularized version of the tophat. Figure 2.1j shows the charge density $k'_\epsilon(x)$ of the regularized elementary kernel, which is the superposition of the charge densities $k'_1(x)$ and $k'_{2,\epsilon}(x)$ in Figs. 2.1k,l.

In principle, other forms of regularization are possible, but the form given here performed well. The regularization requires fewer elementary operations than more complex options. This lends to fast run times and fewer dependencies in the GPU implementation. Note that $k_\epsilon(s) \rightarrow k(s)$ as $\epsilon \rightarrow 0$ when s is not an integer, while the choice of c_ϵ ensures that $k_\epsilon(s)$ is continuously differentiable for all s , with a discontinuity in the second derivative $k''_\epsilon(s)$ of order $\mathcal{O}(\epsilon^2)$ at the half-integers. Since $k_\epsilon(s) = 0$ at the integers, a particle feels no self-force; moreover since $k'_\epsilon(s) \approx 1$ away from the integers and $k_\epsilon(s)$ is zero at the half-integers, a particle induces the correct far-field force. Note that the regularized elementary charge density satisfies $\int_0^1 k'_\epsilon(s) ds = 0$, so the regularization preserves charge neutrality. The regularized electric field kernel is defined by

$$K_\epsilon(x, y) = k_\epsilon\left(\frac{x - y}{L}\right), \quad (2.11)$$

and using this in place of the singular kernel $K(x, y)$ yields a regularized form of the Vlasov-Poisson equations.

2.2.4 Electric field at a particle position

The electric field at a particle position is

$$E(x(x_0, v_0, t), t) = \int_0^L \int_{-\infty}^{\infty} K(x(x_0, v_0, t), y) f(y, v, t) dv dy. \quad (2.12)$$

Then changing variables by the flow map, $(y, u) = (x(y_0, u_0, t), v(y_0, u_0, t))$, yields

$$\begin{aligned} E(x(x_0, v_0, t), t) \\ = \int_0^L \int_{-\infty}^{\infty} K(x(x_0, v_0, t), x(y_0, u_0, t)) f(x(y_0, u_0, t), v(y_0, u_0, t), t) J(y_0, u_0, t) du_0 dy_0 . \end{aligned} \quad (2.13)$$

Using Eq. (2.4) for the invariance of the distribution function and setting the Jacobian $J(x_0, v_0, t) = 1$ since the phase flow is divergence-free, we obtain

$$E(x(x_0, v_0, t), t) = \int_0^L \int_{-\infty}^{\infty} K(x(x_0, v_0, t), x(y_0, u_0, t)) f_0(y_0, u_0) du_0 dy_0 . \quad (2.14)$$

Hence Eq. (2.14) expresses the electric field as a convolution of the electric field kernel with the initial electron distribution.

2.3 Numerical method

This section describes the numerical method comprising the phase-space discretization and adaptive mesh refinement, data structures, electric field quadrature and particle motion, remeshing, interpolation, barycentric Lagrange treecode, and GPU implementation.

2.3.1 Phase-space discretization and adaptive mesh refinement

The computational domain is $[0, L] \times [-v_{max}, v_{max}]$, comprising one period in x -space and an interval in v -space. The domain is discretized by Lagrangian particles and panels, where each panel is a rectangle defined by a 3×3 particle grid as shown in Fig. 2.2a. The panels have a tree structure and, as in previous work in fluid dynamics [59], are adaptively refined; for example in Fig. 2.2, (a) the root panel at level 0 is the entire computational domain, (b) the panels at level 1 were refined in x and v , (c) the panels at level 2 were refined only in v , and (d) two panels at level 3 were refined in x and v . In general, the code performs a user-specified number of uniform panel refinements either in x and v , or only in v , and then applies adaptive mesh refinement (AMR) to the leaf panels to further resolve the electron

distribution function $f(x, v, t)$. For each leaf panel P , the AMR scheme checks the variation in the electron distribution over the nine particles indexed by $i, j = 1 : 3$ comprising the panel to enforce the following refinement criterion,

$$\text{var}(P) = \max_{i,j=1:3} f_{ij} - \min_{i,j=1:3} f_{ij} \leq \epsilon_{\text{var}}, \quad (2.15)$$

where ϵ_{var} is a user-specified tolerance. If the criterion in Eq. (2.15) is satisfied, then the scheme proceeds to check the next leaf panel; otherwise the panel is refined and the resulting subpanels become new leaf panels. The process repeats until the refinement criterion is satisfied for all leaf panels or a user-specified maximum number of AMR levels is reached. The decision on whether to refine in x and v or only in v is problem-dependent and examples will be given in the numerical results section below. At the initial time $t = 0$, the electron distribution in Eq. (2.15) is known analytically, while at later times when remeshing is applied it is computed by interpolation as explained below.




Figure 2.2: phase-space discretization by Lagrangian particles and panels, each panel is a rectangle with a 3×3 particle grid (\bullet), panels have a tree structure and are adaptively refined, (a) level 0, root panel comprising the entire computational domain, (b) level 1, uniformly refined in x and v , (c) level 2, refined only in v , (d) level 3, two panels refined in x and v .

2.3.2 Storage of panel and particle information

The particle information is stored in arrays containing the phase-space coordinate (x_i, v_i) and electron distribution function value f_i of each particle indexed by $i = 1 : N$. The

panel information is organized as a tree with the entire computational domain being the root node and the most refined panels being the leaf nodes. Each panel has 2 or 4 child nodes, depending on whether it was refined in the v direction only or in x and v . The tree is stored in an array of structures, where a panel data structure contains the indices of the parent panel, child panels, and neighbor panels in the panel array, as well as the indices of the panel’s member particles in the particle arrays, and boolean variables indicating whether the panel resides on the boundary of the computational domain. The panel data structure is summarized in Table 2.1.

(a) panel data	(b) data representation
parent panel	1 index in panel array
child panels	4, 2, or 0 indices in panel array
neighbor panels	4 indices in panel array indicating left, right, top, bottom neighbor
member particles	9 indices in particle arrays
domain boundary	4 boolean variables indicating left, right, top, bottom boundary

Table 2.1: Panel data, (a) data associated with each panel, (b) representation of data.

2.3.3 Electric field quadrature and particle motion

Given a phase-space discretization in terms of particles and panels as described above, the electric field integral in Eq. (2.14) is computed as a sum of integrals over the leaf panels, where the integral on each leaf is evaluated by the 2D trapezoid rule in each quadrant of the leaf. This is a $\mathcal{O}(N^2)$ method with the advantage of being spectrally accurate in the case of periodic functions such as we have. The higher-order Simpson’s method was investigated but not found to scale significantly better than the trapezoidal quadrature in the cases we studied. Higher-order methods may still prove to be advantageous and will be the subject of future investigations. When substituted into Eq. (2.3), the trapezoidal discretization yields

a set of ODEs for the motion of the particles in phase-space,

$$x'_i = v_i, \quad v'_i = q^2 \sum_{j=1}^N K_\epsilon(x_i, x_j) f_{0j} \Delta x_{0j} \Delta v_{0j} w_j, \quad x_i(0) = x_{0i}, \quad v_i(0) = v_{0i}, \quad i = 1 : N, \quad (2.16)$$

where $K_\epsilon(x, y)$ is the regularized kernel in Eq. (2.11), $f_{0j} = f_0(x_{0j}, v_{0j})$ is the electron density of a particle, $\Delta x_{0j}, \Delta v_{0j}$ are panel dimensions, and w_j are the trapezoid rule weights (1/16 for panel vertices, 1/8 for edge midpoints, 1/4 for the panel midpoint). In this work the ordinary differential equations (ODEs) are solved using the classical 4th order Runge-Kutta (RK) method. The use of the regularized kernel motivates higher-order methods and the computational gain over second-order time-stepping was clear early on. Earlier FARSIGHT implementations used the second-order standard centered-difference leapfrog method. While both methods show the expected accuracy scalings, gains in run time needed for a given accuracy were observed with the 4th-order RK method over the second-order methods. Higher-order and symplectic time integration are possible subjects of future work.

Discrete analogues of the conserved quantities are obtained by applying the trapezoid rule to the integrals in Eq. (2.5),

$$\begin{aligned} Q_N &= \sum_{i=1}^N f_{0i} \Delta x_{0i} \Delta v_{0i} w_i, & I_N &= \sum_{i=1}^N v_i f_{0i} \Delta x_{0i} \Delta v_{0i} w_i, \\ \mathcal{E}_N &= \frac{1}{2} \sum_{i=1}^N v_i^2 f_{0i} \Delta x_{0i} \Delta v_{0i} w_i + \frac{1}{2} \sum_{\ell=1}^M E_\ell^2 \Delta x_\ell w_\ell, \end{aligned} \quad (2.17)$$

where the sums for Q , I , and the first term of \mathcal{E} are over the particles (x_i, v_i) in phase-space, the sum in the second term of \mathcal{E} is over particles with distinct positions ($x_\ell \neq x_k$ if $\ell \neq k$), and w_i, w_ℓ are the appropriate quadrature weights.

What has been described so far is a Lagrangian particle method for the Vlasov-Poisson system. However, a problem arises; the particles are chosen to resolve the initial electron density, but as time goes on their phase-space distribution is progressively sheared and they

may no longer accurately represent the evolving particle distribution. Figure 2.3 shows this deformation in a Landau damping simulation, where a coarse distribution is used for illustration. Accuracy will always degrade eventually due to deformation, regardless of the resolution. The loss of accuracy can be delayed by increasing the number of particles and panels, but this is only a partial solution since it becomes costly for a long-time calculation. It should be noted that the deformation is not due to time discretization errors in solving the ODEs in Eq. (2.16), but instead shows the limitations of using the same fixed set of Lagrangian particles in long-time calculations. This issue has been noted previously in PIC plasma simulations [53] and also occurs in vortex methods for incompressible fluid flow [63, 62].




Figure 2.3: Landau damping simulation by Eq. (2.16) illustrating how a fixed set of Lagrangian particles and panels is progressively deformed, (a) $t = 0$, (b) $t = 3$, (c) $t = 9$.

To overcome this problem and maintain resolution of the electron phase-space density over long times, we employ remeshing (also called remapping). The benefit of remeshing has been demonstrated in vortex simulations [64] as well as PIC plasma simulations [53, 65].

2.3.4 Remeshing

Figure 2.4 illustrates the remeshing scheme. For simplicity, assume the particles initially lie on a uniform mesh with regular panels as in Fig. 2.4a (in practice they can lie on an adaptively refined mesh as in Fig. 2.2d or Fig. 2.3a). After some time, the particles and panels are deformed as in Fig. 2.4b. The remeshing scheme replaces these particles and

panels by new particles and panels on a new uniform regular mesh, as in Fig. 2.4c. The electron density is transferred from the old particles to the new particles using biquadratic interpolation in the old panels, and AMR is employed to ensure that the new particles resolve the electron density at the current time.




Figure 2.4: Remeshing, (a) regular particles/panels at initial time, (b) irregular particles/panels after evolving for some time, (c) new regular particles/panels (red), old irregular particles/panels (light blue), electron distribution is transferred from old particles to new particles using biquadratic interpolation in old panels, old particles/panels are then deleted and calculation proceeds with new particles/panels.

Given a new particle with coordinates (x^{new}, v^{new}) , the scheme descends the tree of old deformed panels starting from the root and determines which subpanel contains the particle, then which subpanel at the next level contains the particle, and so on until the leaf level is reached. The search has two possible outcomes; if the new particle lies outside the deformed computational domain, then its density is set to zero; otherwise it lies inside a unique old leaf panel P^{old} and then its density is assigned by interpolation, as described below. In checking whether a new particle lies in a given old subpanel, the check is done using the subpanel vertices and assuming the subpanel edges are straight lines. Efficiency is gained for new particles lying on uniform mesh points, where instead of starting from the root, the search employs neighbor panel information from the previous new particle. In problems with periodic boundary conditions, the neighbor panel search takes x -periodicity into account.

2.3.5 Interpolation

The interpolation scheme used in this work is biquadratic. Bilinear interpolation was investigated first but was found to be too diffusive. Using more points per panel and implementing higher-order interpolation schemes could be the subject of future work. Consider a new particle (x^{new}, v^{new}) lying in an old panel P^{old} with its 3×3 grid of particles and their electron density, $(x_{ij}^{old}, v_{ij}^{old}, f_{ij}^{old}), i, j = 1 : 3$. The old panel data defines a biquadratic polynomial,

$$p(x, v) = \sum_{i,j=1}^3 c_{ij} (x - x_{22}^{old})^{i-1} (v - v_{22}^{old})^{j-1}, \quad (2.18)$$

where $(x_{22}^{old}, v_{22}^{old})$ is the panel midpoint particle, and the coefficients c_{ij} are determined by the interpolation conditions,

$$p(x_{ij}^{old}, v_{ij}^{old}) = f_{ij}^{old}, \quad i, j = 1 : 3. \quad (2.19)$$

Then the electron density of the new particle is set to $f^{new} = p(x^{new}, v^{new})$.

There remains the question of how often remeshing should be carried out. It is often thought that remeshing should be done infrequently, to diminish numerical errors arising from interpolation. Hence, in most previous plasma simulations, remeshing was performed every several time steps [1, 53, 65]. In the present work, remeshing is carried out every time step; the rationale is that the phase-space density varies relatively little from one time step to the next, and the interpolation scheme is able to accurately resolve these variations, while a longer interval between remeshings would stress the accuracy of the interpolation scheme. This is a conservative decision and future tuning of the algorithm may find that it is more efficient to remesh less often.

2.3.6 Barycentric Lagrange treecode

The discrete approximation of the electric field in Eq. (2.16) has the form of an N -body interaction,

$$E_i = \sum_{j=1}^N K_\epsilon(x_i, x_j) q_j, \quad q_j = q f_{0j} \Delta x_{0j} \Delta v_{0j} w_j, \quad i = 1 : N, \quad (2.20)$$

where x_i denotes a target particle and x_j are source particles with corresponding charges q_j . Note that the number and charge of the particles changes each at each remeshing. If the field E_i is evaluated by direct summation using loops over i, j , then the operation count is $\mathcal{O}(N^2)$, which is costly for large N . In this work we employ the barycentric Lagrange treecode (BLTC) to reduce the operation count to $\mathcal{O}(N \log N)$ [66, 67]. More particularly, we use the implementation developed in [58].

The BLTC starts by partitioning the particles into clusters, where a cluster in this context is a subinterval of the x -domain $[0, L]$. The set of all clusters has a hierarchical tree structure as depicted in Fig. 2.5a. Given a target particle x_i , the sum over source particles x_j is replaced by a sum over clusters, and these fall into two categories; for clusters C_{near} that are near x_i , the particle-particle interactions are performed directly, while for clusters C_{far} that are far from x_i , the interactions are approximated using barycentric Lagrange interpolation [68]. The resulting approximation has the form,

$$E_i \approx \sum_{C_{\text{near}}} \sum_{x_j \in C_{\text{near}}} K_\epsilon(x_i, x_j) q_j + \sum_{C_{\text{far}}} \sum_{k=0}^n K_\epsilon(x_i, s_k) \hat{q}_k, \quad \hat{q}_k = \sum_{x_j \in C_{\text{far}}} L_k(x_j) q_j, \quad (2.21)$$

where n is the degree of polynomial interpolation, s_k are Chebyshev points adapted to a cluster C_{far} , \hat{q}_k are associated charges, and $L_k(x)$ are the Lagrange interpolating polynomials. A cluster is considered to be far from a target particle when the multipole acceptance criterion (MAC) is satisfied, $r/R < \theta$, where r is the cluster radius, R is the distance between the target and the cluster center, and θ is a user-specified parameter. Figure 2.5b,c depicts the

particle-particle interactions.




Figure 2.5: Treecode schematic, (a) the particles are partitioned into clusters (subintervals of $[0, L]$) with a hierarchical tree structure, in (b,c) the target particles (left, \bullet) interact with (b) source particles (right, \bullet) or with (c) Chebyshev points (right, \times) if there are enough source particles as determined by the MAC, cluster radius r , particle-cluster distance R .

The barycentric form of the Lagrange polynomials is employed,

$$L_k(x) = \frac{w_k}{\sum_{k'=0}^{n-1} \frac{w_{k'}}{x - s_{k'}}}, \quad w_k = (-1)^k \delta_k, \quad \delta_k = \begin{cases} 1/2, & k = 0, n, \\ 1, & k = 1 : n - 1, \end{cases} \quad (2.22)$$

which is computationally stable and efficient [68]. Note that the proxy charges \hat{q}_k of a cluster C in Eq. (2.21) can be precomputed and reused for all targets x_i interacting with C . The operation count of the BLTC is $O(N \log N)$, where the factor N is the number of target particles and the factor $\log N$ is the number of levels in the tree.

2.3.7 GPU implementation

A key feature of the BLTC is that the electric field approximation E_i in Eq. (2.21) consists of direct sums between the target particle x_i and either the source particles $x_j \in C_{\text{near}}$ or the Chebyshev points $s_k \in C_{\text{far}}$, but in either case the necessary kernel evaluations ($K_\epsilon(x_i, x_j), K_\epsilon(x_i, s_k)$) are independent of each other and can be computed concurrently. This enables an efficient GPU implementation of the BLTC using OpenACC [58]. The particle interactions are organized into target batch/source cluster interactions which efficiently map

onto the GPU, where target batching provides an outer level of parallelism and the direct sum form of the particle-particle and particle-cluster interactions provides an inner level of parallelism.

2.3.8 BLTC performance

This section documents the accuracy and efficiency of the BLTC with the regularized kernel developed in section 2.2.3 for the field solver used in FARSIGHT. The computations were done on a single node of the Great Lakes cluster at the University of Michigan, where each node has 36 CPU cores (3.0 GHz Intel Xeon Gold 6154 processors) and one GPU (NVIDIA Tesla V100 with 5120 CUDA cores). The code was written in C++ and compiled with a PGI compiler. The CPU runs were parallelized using OpenMP. For these tests we consider a single evaluation of the electric field in Eq. (2.20), with particles distributed on a cold wave,

$$x_i = i\Delta x + 0.5 \cos(2\pi i\Delta x), \quad \Delta x = 1/N, \quad f_i = 1, \quad i = 1 : N. \quad (2.23)$$

The regularization parameter is set to $\epsilon = 0.1$. The maximum number of particles per leaf is $N_0 = 3000$. We consider the effect of the system size N , MAC parameter θ , and degree n on the BLTC performance.

Figure 2.6 plots the BLTC error versus system size N for a representative value of the MAC parameter $\theta = 0.6$ and several values of the barycentric Lagrange interpolation degree $n = 2 : 2 : 12$. The error is the relative 2-norm of the difference between the field computed by direct summation and by the BLTC, and the results agree to machine precision for calculations on the 36 CPU cores and the GPU. The results show that for a given system size the error decreases as the degree increases, and for a given degree the error increases only slightly as the system size increases. Note that polynomial interpolation at Chebyshev points is spectrally accurate for smooth functions, but the jump in the 2nd derivative of

the regularized electric field constrains the rate of convergence in this case and the error decreases more slowly as the degree increases.

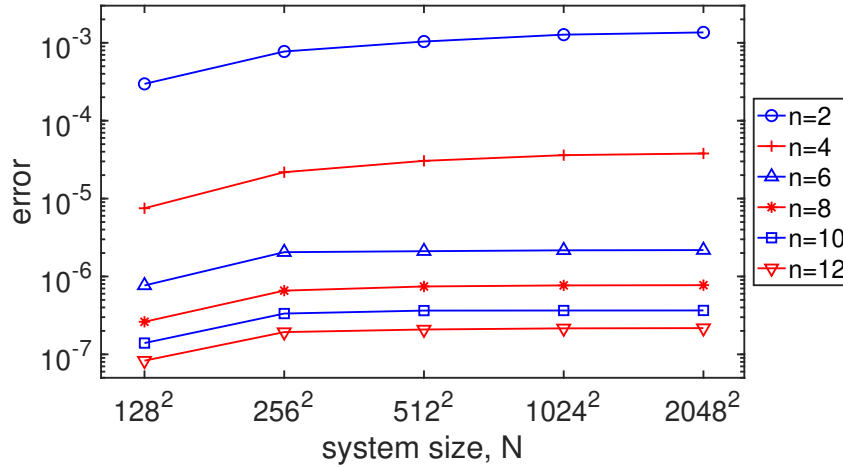


Figure 2.6: BLTC performance, single field evaluation, particles given by Eq. (2.23), error versus system size N , MAC $\theta = 0.6$, degree $n = 2 : 2 : 12$.

Figure 2.7 plots the run time (s) versus the error with four values of the MAC parameter $\theta = 0.8, 0.6, 0.4, 0.2$ constant on each connected curve and several values of the degree $n = 2 : 2 : 12$ given by symbols increasing from right to left on each curve. The calculations use system size $N = 1024^2$. Figure 2.7a shows results for 36 CPU cores and Fig. 2.7b shows results for one GPU. The error decreases as the MAC θ decreases and as the degree n increases. The horizontal line in each plot is the direct sum run time. The BLTC is faster than direct summation and the GPU is faster than 36 CPU cores. These results can be used to select optimal BLTC parameters for a given level of accuracy; for example to achieve about 6-digit accuracy on the GPU, we choose MAC $\theta = 0.6$ and degree $n = 8$. A parameter we don't vary here is the maximum number of particles per leaf; this is set to 3000. It is important to recall that the quadtree structure of panels in phase-space is distinct from the tree used to cluster points for evaluation of the electric field; the former tree has 9 points per leaf panel whereas the treecode structure has a variable number of particles in leaf nodes, with a user-defined maximum number of particles per leaf. This is a tunable parameter that

is set to 3000 for the computations in this dissertation.

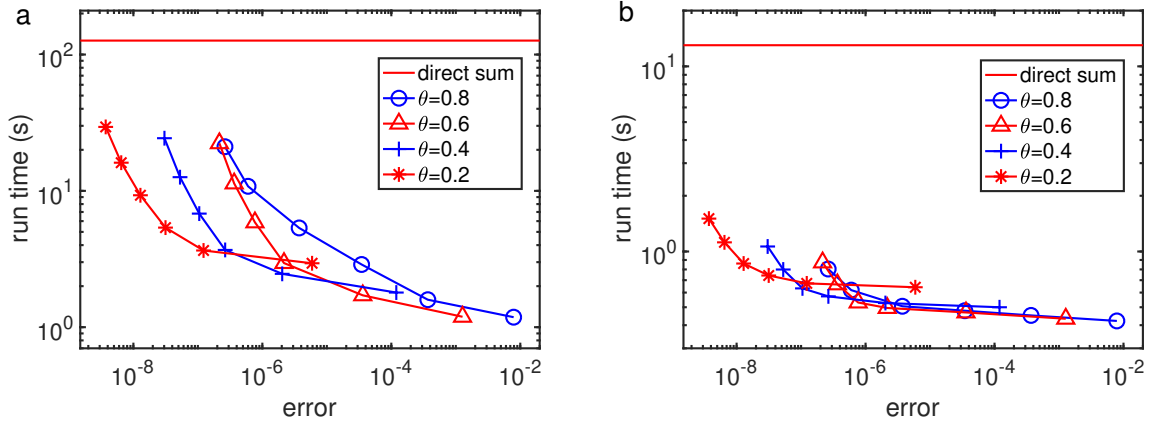


Figure 2.7: BLTC performance, single field evaluation, particles given by Eq. (2.23), run time (s) versus error, $N = 1024^2$ particles, MAC $\theta = 0.2, 0.4, 0.6, 0.8$ constant on each curve, symbols give degree $n = 2, 4, 6, 8, 10, 12$ increasing from right to left on each curve, (a) 36 CPU cores, (b) one GPU, note different vertical scales in (a) and (b), horizontal lines give direct sum run time (s).

Figure 2.8 plots the run time (s) versus system size N for the direct sum and BLTC on 36 CPU cores and one GPU. The BLTC parameters are MAC $\theta = 0.6$ and degree $n = 8$ yielding about 6-digit accuracy. The direct sum run time scales like $O(N^2)$, while the BLTC run time scales like $O(N \log N)$, where the $\log N$ factor varies only slightly over this limited range of N . For system size $N \leq 128^2$, the GPU direct sum is the most efficient option, but for system size $N \geq 512^2$, the GPU BLTC is the most efficient option. For system size $N = 2048^2$, the direct sum run time is 2167 s on 36 CPU cores and 214.8 s on one GPU, while the BLTC run time is 23.6 s on 36 CPU cores and 2.5 s on one GPU. While work was done by Vaughn et al. [58] to ensure the treecode utilizes the GPU efficiently, there is substantial room to improve the performance of the field solver and the treecode on both CPUs with OpenMP and on GPU with OpenACC or another GPU-accessing language. This is a possible avenue of future work.

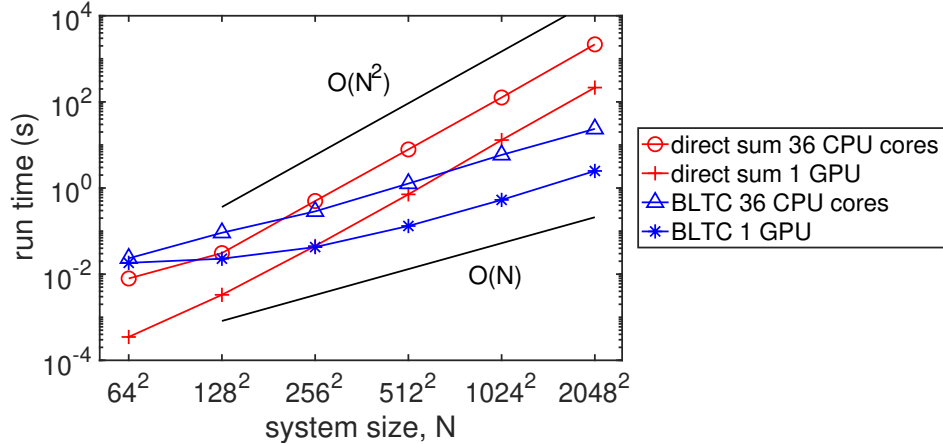


Figure 2.8: BLTC performance, single field evaluation, particles given by Eq. (2.23), MAC $\theta = 0.6$, degree $n = 8$ yielding error $\approx 10^{-6}$, run time (s) vs. system size N , direct sum and BLTC on 36 CPU cores and 1 GPU.

2.4 Numerical Results

The Maxwellian distribution with thermal width v_{th} defined by

$$f_M(v; v_{\text{th}}) = \frac{1}{\sqrt{2\pi}v_{\text{th}}} \exp\left(-\frac{v^2}{2v_{\text{th}}^2}\right), \quad -\infty < v < \infty, \quad (2.24)$$

is a spatially independent equilibrium of the Vlasov-Poisson system. First we consider several problems involving the stability of the Maxwellian with respect to spatial perturbations including Landau damping and two-stream instability. These problems are posed with periodic boundary conditions in x -space and cutoff $v_{\text{max}} = 6$ in v -space. The BLTC parameters are MAC $\theta = 0.6$ and degree $n = 8$, yielding about 6-digit accuracy in the treecode approximation to the discrete convolutions sums for the electric field. Afterwards we consider a problem involving halo formation around an ion beam propagating in free-space. The computations were done on one GPU with maximum number of particles per leaf $N_0 = 3000$. Additional numerical parameters will be specified when each problem is described.

2.4.1 Landau damping

In this case the initial electron density is a spatially perturbed Maxwellian,

$$f_0(x, v) = f_M(v; v_{\text{th}})(1 + \alpha \cos(kx)), \quad (x, v) \in [0, 4\pi] \times [-6, 6], \quad (2.25)$$

with perturbation amplitude α , wavenumber $k = 0.5$, and thermal width $v_{\text{th}} = 1$.

Weak Landau damping

First we consider small perturbation amplitude, $\alpha = 0.01$. Figure 2.9 plots the 2-norm of the electric field versus time for phase-space discretization $N = 256^2$, time step $\Delta t = 0.25$, and regularization parameter $\epsilon = 0.1$. The electric field decays exponentially due to Landau damping, as originally derived for the linearized Vlasov-Poisson equations [69], and the computed results are in good agreement with the predicted decay rate $\gamma = 0.1533$ [1] indicated by the dashed line.




Figure 2.9: Weak Landau damping, 2-norm of electric field versus time, phase-space discretization $N = 256^2$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, dashed line gives predicted decay rate $\gamma = 0.1533$ [1], GPU run time 34.6 s.

Figure 2.10 shows the dependence of the electric field on the numerical parameters. The top row varies the phase-space discretization N , the middle row varies the time step Δt , and the bottom row varies the regularization parameter ϵ . The computed electric field converges to the correct decay rate as the numerical parameters are refined.

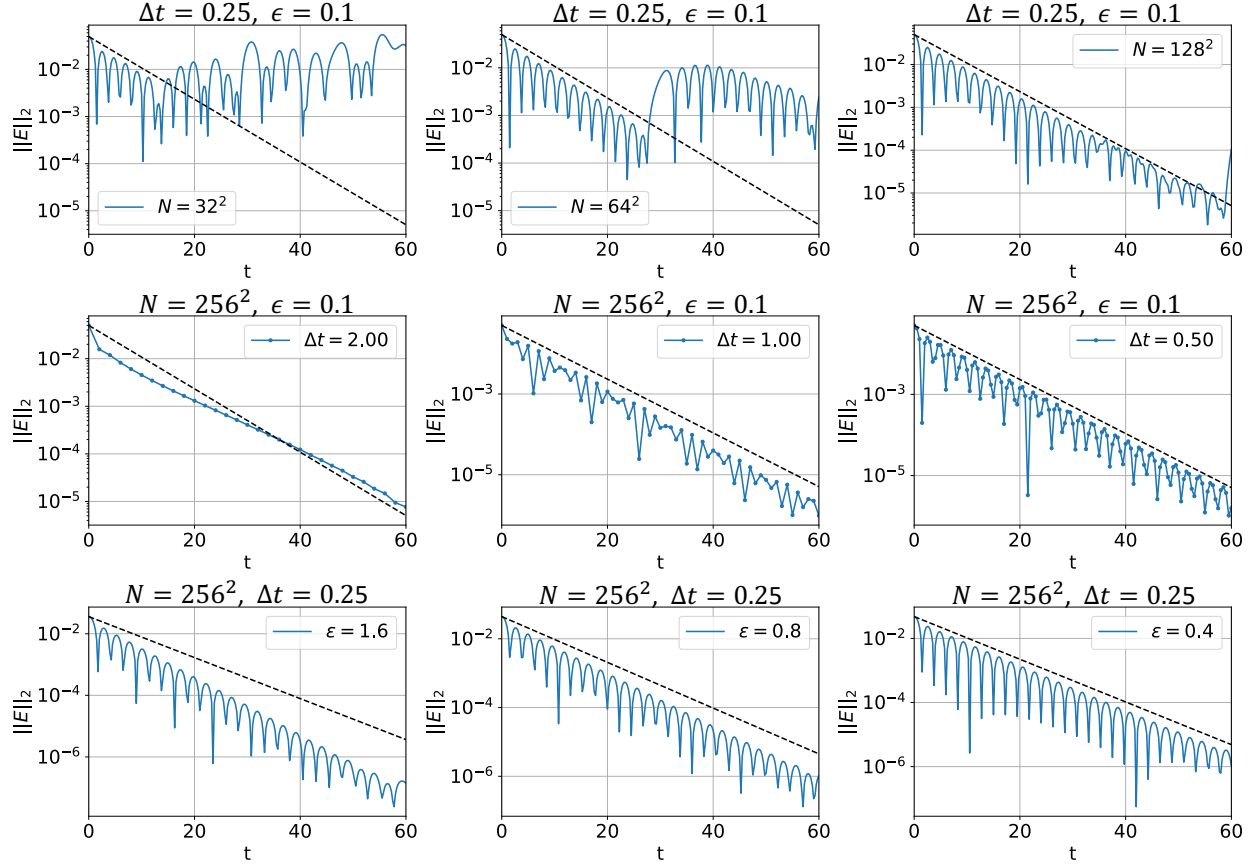


Figure 2.10: Weak Landau damping, 2-norm of electric field versus time, effect of numerical parameters, top row: phase-space discretization $N = 32^2, 64^2, 128^2$, middle row: time step $\Delta t = 2, 1, 0.5$, bottom row: regularization parameter $\epsilon = 1.6, 0.8, 0.4$.

Next we consider the electron distribution. Figure 2.11a shows the initial distribution, where the perturbation is too small to be seen. Figures 2.11b-f plot the perturbation from the Maxwellian, $\delta f(x, v, t) = f(x, v, t) - f_M(v)$, where blue indicates an excess and red indicates a deficit. At time $t = 0$ the perturbation is an alternating sequence of blue and red patches located symmetrically around velocity $v = 0$. Due to free-streaming, portions of a patch with $v > 0$ travel to the right and portions with $v < 0$ travel to the left, but in addition some electrons accelerate to higher velocity. The patches are progressively stretched, but their width is not uniform. The patches enter neighboring spatial periods and overlap in physical space, tending to a stack of nearly horizontal thin filaments. To resolve this structure,

this computation used phase-space discretization $N = 1024^2$, time step $\Delta t = 0.125$, and regularization parameter $\epsilon = 0.05$.

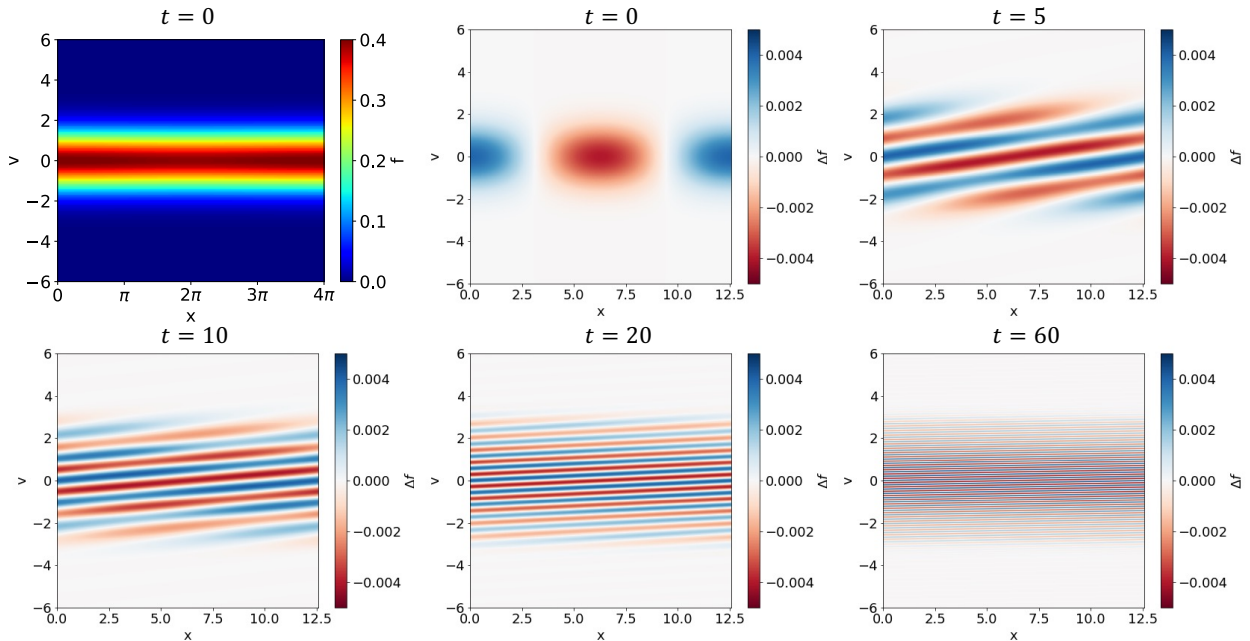


Figure 2.11: Weak Landau damping, (a) initial electron distribution $f_0(x, v)$, (b)-(f) perturbation from Maxwellian, $\delta f(x, v, t) = f(x, v, t) - f_M(v)$, time $t \in [0, 60]$, phase-space discretization $N = 1024^2$, time step $\Delta t = 0.125$, regularization parameter $\epsilon = 0.05$, run time 1372 s.

Figure 2.12 plots slices of the electron density perturbation from the Maxwellian, $\delta f(0, v, t) = f(0, v, t) - f_M(v)$, at a representative spatial location $x = 0$. As time proceeds, the perturbation oscillates more rapidly and the envelope broadens into a roughly triangular shape. This gives a physical interpretation of Landau damping; the electron distribution $f(x, v, t)$ tends to zero weakly, while the charge density $\rho(x, t)$, which integrates $f(x, v, t)$ over v , tends to zero pointwise, and this in turn causes the electric field to decay. The weak convergence seen here has been rigorously proven for the full nonlinear Vlasov-Poisson equations [70].

Next we demonstrate convergence of the computed electron distribution at time $t = 60$ for the small amplitude case. The reference is a highly resolved calculation with phase-space discretization $N = 2048^2$, time step $\Delta t = 0.125$ and regularization parameter $\epsilon = 0.05$. In

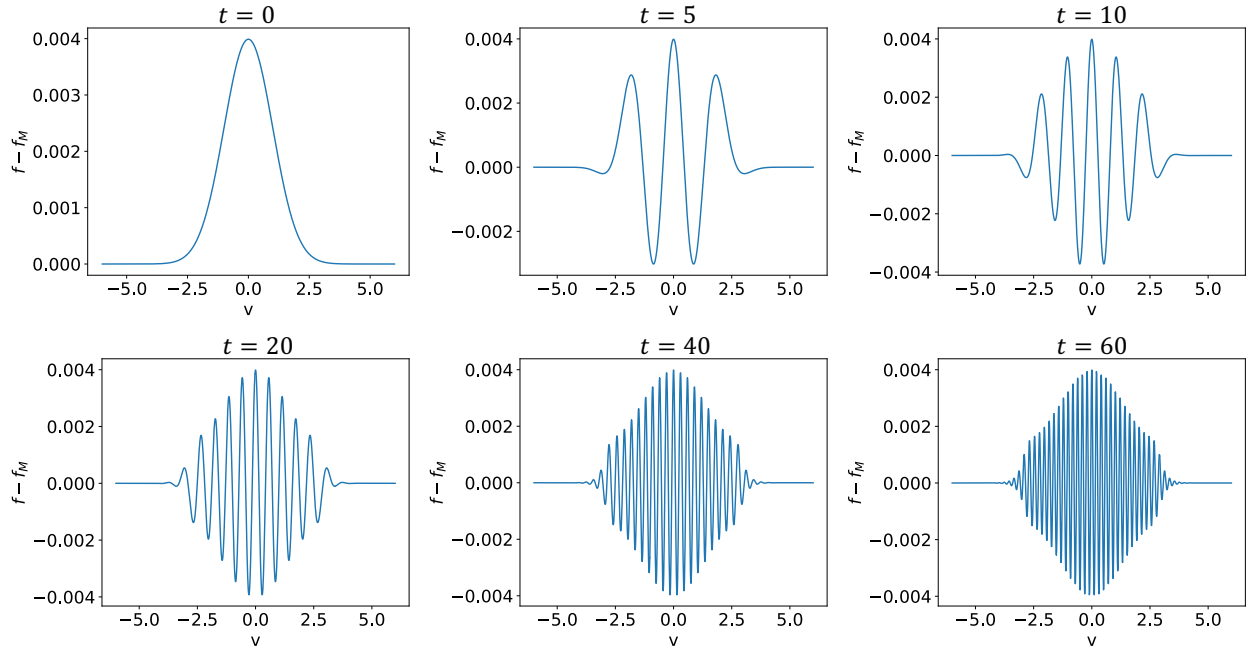


Figure 2.12: Weak Landau damping, slices through electron density deviation from Maxwellian, $\delta f(0, v, t) = f(0, v, t) - f_M(v)$, time $t \in [0, 60]$, same parameters as Fig. 2.11.

each case, one parameter was varied and the other two were kept fixed. Figure 2.13a shows that the error scales like $O(\Delta x^3)$, the expected rate for biquadratic interpolation and 2D trapezoid quadrature. Figure 2.13b shows that the error scales like $O(\Delta t^4)$, the expected rate for 4th order Runge-Kutta time-stepping. Figure 2.13c shows that the error scales like $O(\epsilon^2)$, which is the rate at which the regularized electric field kernel in Eq. (2.10) converges away from the singularity. Convergence with respect to a regularization parameter has been considered in analyses of particle methods [71, 62, 72], but we are not aware of any rigorous result for the present case.

Next we consider the conserved quantities in Eqs. (2.17). Fig. 2.14 presents a sequence of calculations with refinements in the regularization parameter ϵ , number of particles N , and time step Δt , showing improved conservation as the numerical parameters are refined. In Fig. 2.14 the variations in (a) total charge Q_N , (b) total current I_N , and (c) total energy \mathcal{E}_N are shown. In the case of weak Landau damping, the minimal distortion of the phase-space




Figure 2.13: Weak Landau damping, convergence of electron density with respect to numerical parameters, time $t = 60$, reference solution computed with phase-space discretization $N = 2048^2$, time step $\Delta t = 0.125$, regularization parameter $\epsilon = 0.05$, solid line shows asymptotic error scaling, (a) $N \in [32^2, 1024^2]$, error = $O(\Delta x^3)$, (b) $\Delta t \in [0.25, 2]$, error = $O(\Delta t^4)$, (c) $\epsilon \in [0.1, 0.8]$, error = $O(\epsilon^2)$.

panels implies that the remeshing interpolation is highly accurate, and each refinement in numerical parameters gives an extra order of magnitude better conservation in total initial charge Q_{N0} and total initial current I_{N0} and at least $2\times$ better conservation of total initial energy \mathcal{E}_{N0} .




Figure 2.14: Weak Landau damping, variation in conserved quantities vs. time, effect of numerical parameters, (a) charge Q_N , (b) current I_N , (c) energy \mathcal{E}_N , phase-space discretization N , time step Δt , regularization parameter ϵ , blue (—): $(N, \Delta t, \epsilon) = (128^2, 1.0, 0.4)$, orange (—): $(N, \Delta t, \epsilon) = (256^2, 0.5, 0.2)$, green (—): $(N, \Delta t, \epsilon) = (512^2, 0.25, 0.1)$.

Strong Landau damping

Figure 2.15 plots the 2-norm of the electric field versus time with large perturbation amplitude, $\alpha = 0.5$. In this case the electric field decays until time $t = 10$, then grows until time $t = 40$, and then oscillates in time. The damping rate $\gamma_1 = 0.2920$ and growth rate $\gamma_2 = 0.0815$ are illustrated by the dashed lines in Fig. 2.15 and documented in the

literature [2, 73, 74].




Figure 2.15: Strong Landau damping, 2-norm of electric field vs. time, phase-space discretization $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, run time 754 s, decay rate ($\gamma = -0.2920$, orange line), growth rate ($\gamma = 0.0815$, black line) [2].

Figure 2.16 shows the corresponding electron distribution, which exhibits significant non-linear particle-wave interactions. The growth of the electric field starting at time $t = 10$, is correlated with particle trapping as indicated by the formation of phase-space vortices centered at the wave phase velocity, $v_{ph} = \omega_r/k \approx 2.75$, since $k = 1/2$ and ω_r is approximately determined by the Bohm-Gross dispersion relation, $\omega_r \approx 1 + \frac{3}{2}(kv_{th})^2 = 11/8$. The vortices develop gradually in time and are clearly visible by time $t = 36$. The vortices show complex evolution between times $t = 36$ and $t = 60$, traveling at approximately the estimated phase velocity.

Figure 2.17 shows the behavior of conservation metrics with refinement in parameters. Fig. 2.17(a) shows the variation in total numerical charge Q_N , (b) total current I_N , (c) total energy \mathcal{E}_N , (d) f_{\max} , (e) f_{\min} , and (f) fraction of negative distribution f , $F_{N-} \equiv |Q_{N-}/Q_N|$, where the integrated negative f is $Q_{N-} \equiv q \sum_{i=1}^N f_{0i} \Delta x_{0i} \Delta v_{0i} w_i$. The fraction of negative f in Fig. 2.17(f) is included because FARSIGHT does not employ any limiters to enforce positivity or mass conservation. This diagnostic provides a measure of the interpolation error. We note that all errors are controlled as numerical parameters are refined. Also, for the fraction of negative f and the maximum and minimum f values, the errors start




Figure 2.16: Strong Landau damping, electron distribution $f(x, v, t)$, discretization $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, run time 754 s.

at one level and increase at some time where nonlinear mixing stresses the simulation, and this mixing time also increases with a refinement in parameters. The conservation in energy is comparable with the literature [1, 75, 76].

The calculations presented so far used a uniform mesh. In Fig. 2.18 we present the AMR phase-space for the strong Landau problem. In this computation, which took 345 s, we use $N = 16 \times 64 + 5$ levels of refinement and the variation criterion with $\epsilon_{AMR} = 0.01$. The dynamic adaptive refinement tracks the filaments and vortices as they evolve, demonstrating the ability of the AMR scheme to efficiently resolve features in phase-space. The panels at




Figure 2.17: Strong Landau damping, conservation properties, BLTC parameters MAC $\theta = 0.6$, degree $n = 8$, mesh size N , regularization parameter ϵ , time step Δt refined simultaneously, (a) total charge Q_N , (b) total current I_N , (c) total energy \mathcal{E}_N , (d) f_{\max} , (e) f_{\min} , (f) fraction of negative f , $|Q_{N-}/Q_N|$, blue (—): $(N, \Delta t, \epsilon) = (128^2, 0.5, 0.8)$, orange (—): $(N, \Delta t, \epsilon) = (256^2, 0.25, 0.4)$, green (—): $(N, \Delta t, \epsilon) = (512^2, 0.125, 0.2)$, GPU times (s) 19.3, 156.9, 1356.4

the highest level are too small to be seen clearly in the figure, so instead they are shaded by level of refinement. In Fig. 2.19 a smaller region of phase-space at time $t = 60$ is shown so that the panels can be seen. Fig. 2.19(a) shows the phase space at time $t = 60$ with a red box indicating the smaller region that is looked at in greater detail in (b). Fig. 2.19(c) shows the actual panel outlines as opposed to a heat map coloring panels by AMR level. This affords a different perspective on the ability of the AMR scheme to resolve phase-space features, and illustrates the amount of detail attained with the highest level of refinement.

The uniform and AMR phase-spaces are nearly indistinguishable, as seen by comparing the uniform refinement calculation performed with resolution $N = 512 \times 2048$, shown in Fig. 2.16, and the AMR calculation performed with $N = 16 \times 64 + 5$ levels of refinement, shown in Fig. 2.18. The uniform resolution computation at $N = 512 \times 2048$ took 754 s whereas the AMR computation took 345 s; the AMR computation is more than $2\times$ faster than the uniform calculation. The AMR calculation began with $N = 154297$ points which

is 14.7% of the number of points in the uniform calculation, and ended with $N = 538601$ which is 51.2% of the uniform N .




Figure 2.18: Strong Landau damping, AMR simulation, $t_f = 60$, $N = 16 \times 64 + 5$ AMR levels, $N_0 = 154297$, $N_f = 538601$, or 14.7% to 51.2% of the number of points in the uniform calculation, comparable uniform resolution is $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, $\epsilon_{AMR} = 0.01$, GPU time 345 s, top: phase-space, bottom: panel level.

2.4.2 Two-stream instability

There are several commonly studied forms of the two-stream instability. Following [1, 46], we consider the initial distribution

$$f_0(x, v) = v^2 f_M(v; v_{th}) (1 + \alpha \cos(kx)), \quad (x, v) \in [0, 4\pi] \times [-6, 6], \quad (2.26)$$

with thermal width $v_{th} = 1$, perturbation amplitude $\alpha = 0.5$ and wavenumber $k = 0.5$. Computations are carried out to time $t = 60$ with cutoff velocity $v_m = 6$, phase-space discretization $N = 16 \times 64 + 5$ AMR levels, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, and AMR tolerance $\epsilon_{AMR} = 0.008$. Figure 2.20 shows the electron phase-space density. The first row shows times 0, 3, 6, and 9 highlighting the formation of the central vortex that is characteristic of the two-stream instability. The second row shows times




Figure 2.19: Strong Landau damping, AMR simulation, $N = 16 \times 64 + 5$ AMR levels, comparable uniform resolution is $N = 512 \times 2048$, regularization parameter $\epsilon = 0.1$, time step $\Delta t = 0.25$, AMR tolerance $\epsilon_{AMR} = 0.01$, run time 345 s, (a) phase-space at time $t = 60$, (b) zoom of small region in phase-space centered at $(x, v) = (2, 2)$, (c) phase-space panels in zoom region of (b).

between 12, 15, 18, and 20. Here we see the instability of the cat's eye at $x = 0/4\pi$ as phase-space density crosses between periods. The third row shows times 49, 50.5, 51.5, and 60 and the formation of secondary vortices on the outer tails of the distribution.

While the central phase-space vortex of the two-stream instability is not present in the Landau damping scenario, similarities can be observed between the two problems. At high velocities, the entrainment of low density phase fluid forms small vortices near velocity $v = \pm 3$ and these small vortices propagate from one period to the next. As in the strong Landau damping simulations, thin filaments develop as the electron distribution is stretched and wrapped around the phase-space vortices.

The AMR computation with phase-space discretization $N = 16 \times 64 + 5$ AMR levels took 437 s and the corresponding uniform calculation with $N = 512 \times 2048$ took 745 s. The phase-spaces for the two computations at time $t = 60$ are compared in Fig. 2.21. Fig. 2.21(a) shows the uniform phase-space calculation at time $t = 60$, (b) shows the AMR phase-space, and (c) shows the level of panel refinement. As in the strong Landau damping problem, the uniform and AMR phase-spaces are nearly indistinguishable and the image of refinement level indicates the ability of the AMR to track features in phase-space. In the two-stream

instability, features develop throughout most of phase-space and so most of phase-space is maximally resolved. The next examples will have features more localized in phase-space and thus highlight further the benefit of AMR.




Figure 2.20: Two-stream instability, phase-space electron distribution, initial condition Eq. (2.26), discretization $N = 16 \times 64 + 5$ AMR levels, $N = 264325$ initially, which is 25.1% of the uniform N and $N = 639165$ ultimately which is 60.8% of the uniform N , regularization parameter $\epsilon = 0.1$, run time 437 s compared to uniform run time 739 s.




Figure 2.21: Two-stream instability, time $t = 60$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, compare AMR with uniform resolution, (a) uniform discretization, $N = 528 \times 2048$, run time 739 s, (b) AMR, $\epsilon_{AMR} = 0.008$, $N = 16 \times 64 + 5$ AMR levels, equivalent to 25.1% to 60.8% of uniform N , run time 435 s, (c) AMR panels shaded by refinement level.

2.4.3 Cold two-stream instability

Next we consider a problem in which the initial density has two Maxwellians centered at $v = \pm v_s$ with a spatially periodic perturbation,

$$f_0(x, v) = \frac{1}{2} (f_M(v - v_s; v_{th}) + f_M(v + v_s; v_{th})) (1 + \alpha \cos kx), \quad (x, v) \in [0, 4\pi] \times [-6, 6]. \quad (2.27)$$

In these computations the thermal width is $v_{th} = 0.25$ and the stream velocity is $v_s = 2$, yielding two well-separated cold streams. The perturbation amplitude is $\alpha = 0.5$ and the wavenumber is $k = 0.5$.

We pursued this problem in order to test the AMR on a problem that would be especially costly to resolve uniformly. The resulting phase-spaces resemble problems of practical interest in plasma-based acceleration, where particles are accelerated to such high energies that the background plasma has a relatively small thermal energy. The cold two-stream instability phase-spaces are shown in Fig. 2.22. Times $t = 0$ to $t = 10$ show acceleration of some particles to velocities $2.5\times$ the initial streaming velocity as filaments are pushed to $v = \pm 5$. By time $t = 16$ the central phase-space vortex is well established and there is

some crossing of phase-space density between periods. Between times $t = 14$ and $t = 22$ the faint high-velocity filaments frame secondary vortices on the higher velocity slopes of the two initial streams as in the previous warm two-stream instability.

The calculation in Fig. 2.22 has discretization is $N = 16 \times 32 + 5$ levels of AMR and uses AMR tolerance $\epsilon_{AMR} = 0.008$. The uniform and AMR calculations are compared in Fig. 2.23. Fig. 2.23(a) shows the uniform phase-space with phase-space discretization $N = 512 \times 1024$, (b) shows the AMR phase-space with discretization $N = 16 \times 32 + 5$ levels of AMR, and (c) shows the refinement levels in the AMR calculation. In the AMR calculation, the initial N is 129245, which is 24.6% of the uniform N , and the final N is 193345, which is 36.8% of the uniform resolution. The uniform resolution simulation took 123.5 seconds and the AMR calculation took 44.1 seconds; the uniform calculation took 2.8 times longer.

2.4.4 Halo formation in a mismatched thermal equilibrium sheet beam

Our last example concerns a thermal equilibrium sheet beam previously simulated by Campos Pinto et al. using the LTPIC method [4]. The model describes an axially thin transverse slice of a continuously focused ion beam in an accelerator [3]. The ion number distribution of the slice is $f(x, x', s)$, where x is a transverse spatial coordinate, x' is the angle between particle trajectories and the longitudinal machine axis, and s is the timelike axial coordinate of a reference particle along the design orbit in the machine. The phase-space (x, x') is infinite, $-\infty < x, x' < \infty$. The ion distribution is given at $s = 0$ and the goal is to compute $f(x, x', s)$ for $s > 0$.

In dimensional form, the Vlasov-Poisson equations describing the transverse phase-space are

$$\partial_s f + x' \partial_x f + \left(\frac{q}{m \gamma_b^3 \beta_b^2 c^2} E - \kappa x \right) \partial_{x'} f = 0, \quad E = -\partial_x \phi, \quad -\partial_x^2 \phi = \frac{q}{\epsilon_0} n, \quad (2.28)$$

where $E(s, x)$ is the electric field, $\phi(s, x)$ is the potential, $n(s, x) = \int_{-\infty}^{\infty} f(x, x', s) dx'$ is the




Figure 2.22: Cold two-stream instability, initial distribution Eq. (2.27), phase-space discretization $N = 16 \times 32 + 5$ AMR levels, $N_0 = 129245$ which is 24.6% of the uniform N , N_f is 193345, which is 36.8% of the uniform N , AMR tolerance $\epsilon_{AMR} = 0.008$, regularization parameter $\epsilon = 0.1$, time step $\Delta t = 0.25$, run time 44.1 s.

particle number density, q, m are the ion charge and rest mass, $\gamma_b = (1 - \beta_b^2)^{-1/2}$ is the Lorentz factor, $\beta_b c$ is the beam slice speed, and $\kappa = (2\pi/3)^2 m^{-2} > 0$ is the focusing strength. The potential satisfies the following free-space boundary conditions,

$$-\partial_x \phi(s, x) \rightarrow \pm \frac{q}{2\epsilon_0} N_i, \quad x \rightarrow \pm\infty, \quad (2.29)$$

where $N_i = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(s, x, x') dx dx'$ is the total number of particles in the beam slice.

Lund et al. [3] showed that the model gives rise to an equilibrium distribution,

$$f^{eq}(x, x') = \frac{\hat{n}}{\sqrt{2\pi T^*}} \exp\left(-\frac{x'^2}{2T^*}\right) \tilde{n}_\Delta \left(\frac{x}{\gamma_b \lambda_D}\right), \quad (2.30)$$

where \hat{n} is the peak density, T^* is the dimensionless temperature, λ_D is the Debye length,




Figure 2.23: Cold two-stream instability, compare AMR with uniform resolution at time 22, $\epsilon = 0.1, \Delta t = 0.25$ (a) $N = 512 \times 1024$, GPU time 123.5 s, (b) $N = 16 \times 32 + 5$ levels of AMR, $\epsilon_{AMR} = 0.008$, GPU time 44.1 s, (c) AMR resolution.

$\hat{\omega}_p$ is the plasma frequency, Δ is a dimensionless parameter,

$$\hat{n} = \frac{\epsilon_0 m \gamma_b^3 \beta_b^2 c^2 \kappa}{q^2 (1 + \Delta)}, \quad T^* = \frac{T}{m \gamma_b \beta_b^2 c^2}, \quad \lambda_D = \left(\frac{T}{m \hat{\omega}_p^2} \right)^{1/2}, \quad (2.31)$$

$$\hat{\omega}_p = \left(\frac{q^2 \hat{n}}{\epsilon_0 m} \right)^{1/2}, \quad \Delta = \frac{\gamma_b^3 \beta_b^2 c^2 \kappa}{\hat{\omega}_p^2} - 1,$$

and the scaled density \tilde{n}_Δ satisfies the following equation and normalizing conditions,

$$\tilde{n}_\Delta'' + (1 + \Delta) \tilde{n}_\Delta - \tilde{n}_\Delta^2 + \frac{1}{\tilde{n}_\Delta} (\tilde{n}_\Delta')^2 = 0, \quad \tilde{n}_\Delta(0) = 1, \quad \tilde{n}_\Delta'(0) = 0. \quad (2.32)$$

We solved Eq. (2.32) numerically using the parameter values in Table 2.2 [4], and Fig. 2.24 shows the resulting equilibrium distribution which has a smooth top-hat profile. The connection of the scaled density \tilde{n}_Δ to quantities from Campos Pinto et al. [4] is the relation $\tilde{n}_\Delta = e^{-\psi_\Delta}$ where ψ_Δ is an effective potential defined to be the solution of the transformed Poisson equation,

$$\psi_\Delta''(\hat{x}) = 1 + \Delta - e^{-\psi_\Delta(\hat{x})}, \quad (2.33)$$

satisfying $\psi_\Delta'(0) = \psi_\Delta(0) = 0$. Identifying $\tilde{n}_\Delta = e^{-\psi_\Delta}$, differentiating, and substituting Eqn. (2.33) leads to the Eqn. (2.32) for \tilde{n}_Δ .

At this point it is convenient to scale variables,

$$x \rightarrow \kappa^{1/2} x, \quad \phi \rightarrow \frac{q}{m \gamma_b^3 \beta_b^2 c^2} \phi, \quad f \rightarrow \frac{q^2 \kappa}{m \gamma_b^3 \beta_b^2 c^2 \epsilon_0} f, \quad (2.34)$$

parameter Δ	temperature T^*	Debye length λ_D	peak density \hat{n}
5.522×10^{-8}	3.463×10^{-7}	2.810×10^{-4} m	4.848×10^{13}

Table 2.2: Parameters for equilibrium sheet beam distribution from Campos Pinto et al. [4].

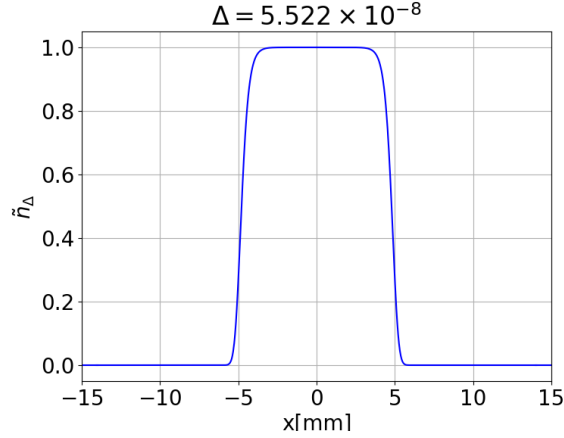


Figure 2.24: Scaled density \tilde{n}_Δ for equilibrium sheet beam distribution [3, 4].

after which the Vlasov-Poisson equations are

$$\partial_s f + x' \partial_x f + (E - x) \partial_{x'} f = 0, \quad E = -\partial_x \phi, \quad -\partial_x^2 \phi = n, \quad n = \int_{-\infty}^{\infty} f dx', \quad (2.35)$$

satisfying open boundary conditions. Following [4] we consider a mismatched beam in which the equilibrium is perturbed,

$$f(x, x', s = 0) = f^{eq}(x/\mu, \mu x'), \quad \mu = 1.25, \quad (2.36)$$

and the resulting ion distribution $f(x, x', s)$ for $s > 0$ is obtained using the FARSIGHT method. In this case the electric field convolution integral is

$$E(s, x) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} K_\epsilon(x, y) f(y, x', s) dy dx', \quad K_\epsilon(x, y) = \frac{1}{2} \frac{x - y}{((x - y)^2 + \epsilon^2)^{1/2}}, \quad (2.37)$$

where the regularized 1D free-space kernel converges to the exact kernel, $K_\epsilon(x, y) \rightarrow \frac{1}{2} \text{sign}(x - y)$, as $\epsilon \rightarrow 0$.

The calculation used phase-space domain $|x| \leq \sqrt{\kappa} 15 \text{ mm} = \pi/100$ and $|x'| \leq 14.5 \text{ mrad}$,

uniform discretization $N = 1024^2$, axial increment $\Delta s = \sqrt{\kappa} L_p/10 = \pi/3$, regularization parameter $\epsilon = \sqrt{\kappa} 0.1 \text{ mm} \approx 2.09 \times 10^{-4}$, and BLTC pararameters MAC $\theta = 0.8$ and degree $n = 4$. Figure 2.25 shows the evolution of the phase-space density with halo formation. The filled colored contours show the density in the core of the beam and the open black contours show the low density halo that emerges around the core (contour levels 10^{-p} , $p = 1 : 5$). The results agree well with the LTPIC reference solution in [4], which used 512^2 linearly transformed (LTPIC) particles each with four degrees of freedom, and $\Delta s = \sqrt{\kappa} L_p/16 = \pi/48$.

Figure 2.26 presents the solution at axial position $s = \sqrt{\kappa} 40L_p = 40\pi/3$ with (a) uniform discretization $N = 1024^4$, (b) discretization $N = 16^2 + 6$ AMR levels, initially $N = 135253$ or 12.9% of N_u , finally $N = 310901$ or 29.6% of N_u and (c) the panels in the AMR calculation. Comparing frames (a) and (b) shows that the uniform and AMR phase-space distributions are nearly indistinguishable. The run time of the uniform calculation was 1169 s, while the run time of the AMR calculation was 245 s. The result in frame (c) shows that the AMR scheme refined panels within the core of the beam and around the halo, leaving the complementary low-density regions unrefined.

Table 2.3 presents some profile information, including the number of initial, final, and average number of particles used as well as the total run time, field evaluation time, and remeshing time for the uniform and AMR calculations. The results show that the AMR calculation is about $5\times$ faster, using on average $5\times$ fewer points. The AMR calculation requires only 15% of the uniform field evaluation time and one third of the uniform remeshing time.




Figure 2.25: Mismatched beam problem, uniform phase-space discretization $N = 1024^2$, axial increment $\Delta s = 0.1L_p$, regularization parameter $\epsilon = 0.1$ mm, GPU run time 1169 s, results agree well with LTPIC reference solution in [4].

2.5 Summary

A forward semi-Lagrangian scheme called FARSIGHT was presented for collisionless electrostatic plasmas described by the 1D1V Vlasov-Poisson equations. The distribution




Figure 2.26: Mismatched beam problem, phase-space at axial position $s = 40L_p$, (a) uniform discretization, $N = 1024^2$, run time 1169 s, (b) AMR, $N = 64^2 + 6$ AMR levels, initially $N = 135253$ or 12.9% of N_u , finally $N = 207721$ or 19.8% of N_u , run time 245 s, (c) panels in AMR calculation.

N	N_0	N_{final}	N_{ave}	total run time	field evaluation time	remeshing time
uniform 1024^2	1050625	1050625	1050625	1169 s	735 s	349 s
$8^2 + 6$ AMR levels	135253	310901	198624	247 s	114 s	116 s
% of uniform	12.9%	29.6%	18.9%	21.1%	15.5%	33.3%

Table 2.3: Mismatched beam profile, profile of results in Fig. 2.26, discretization N , initial and final values N_0 , N_f , average value N_{ave} , total run time, field evaluation time, remeshing time

function is represented by quadrilateral panels having a hierarchical tree structure, and each panel is a 3×3 grid of particles which are tracked by 4th order Runge-Kutta timestepping. The electric field is expressed as a convolution integral of the charge density with a regularized electric field kernel. The particles are remeshed at every time step using biquadratic interpolation on the panels, and the panels are adaptively refined to resolve the phase-space distribution function. The electric field integral is discretized by the trapezoid rule, and the discrete sums are computed by a GPU-accelerated barycentric Lagrange treecode [58].

In calculations of Landau damping, FARSIGHT reproduced the well-known damping and growth rates, while in strong Landau damping, two-stream instability, and beam halo formation, the code was able to resolve small-scale structures in phase-space. The benefit of adaptive mesh refinement was shown by significantly decreasing the run time for the cold

two-stream instability and beam halo formation problems. It was shown that conservation errors decrease and fewer negative f -values occur when the numerical parameters are refined.

FARSIGHT resembles other forward semi-Lagrangian schemes such as LTPIC [4], SFK [31], and the cloudy Vlasov solver [51] in tracking the deformation of the phase-space flow map, however here this is done using Lagrangian panels with a tree structure. The scheme does not use operator splitting.

There are several directions for future development of FARSIGHT. The present simulations did not use a positivity-preserving limiter, but this could be considered in the future. It would be interesting to test the effect of using finer particle grids in each panel and higher order interpolation and quadrature. The adaptive refinement procedure could be improved by enabling panel refinement independently in the x and v directions. We believe the approach developed here can be extended to higher-dimensional electromagnetic plasmas in confined geometry, using an integral formulation of Maxwell's equations, adaptive refinement of phase-space, and a fast summation method that scales well on multiple GPUs [77].

CHAPTER III

Unlimited Photon Acceleration

This chapter presents unlimited photon acceleration (PA^∞), a scheme for unlimited frequency shift of a laser pulse trailing an electron beam through plasma with a tapered density profile.

The many applications of bright, coherent XUV light have motivated substantial interest in source development, such as the construction of XUV wavelength Free Electron Lasers (FELs) including FLASH [78] as well as nonlinear frequency mixing [79], high harmonic generation [80], and XUV lasing [81], to name a few. Another method for generating short wavelengths is ‘photon acceleration’ [8, 16]. Dressed by a background plasma, ‘quasi-photons’ gain an effective mass $\hbar\omega_p/c^2$, where $\omega_p^2 = e^2n/m_e\varepsilon_0$ for a plasma of number density n . In the presence of co-propagating density gradients, the quasi-photons experience local frequency shifts due to spatiotemporal variations in the phase velocity, and are therefore accelerated (i.e. experience an increase in group velocity). The resulting quasi-photon phase-space trajectories in plasma wakefields are similar to those of leptons [82].

Photon acceleration can arise as a result of plasma wakefields [8], ionization fronts [83, 84] and even using metamaterials [85]. Photon acceleration was measured in ionization front [86] and laser wakefield acceleration experiments [87, 88]. Recent results include cascaded sequences of localized ionizations [20], resulting in large frequency shifts and the use of plasma

wakes to downshift radiation to very long wavelengths [22]. Limits to photon acceleration in plasma wakefields in the linear [17] and nonlinear regimes [18] were previously studied, identifying dephasing of the photon beam with respect to the accelerating refractive index gradient as placing a ceiling on the frequency shift. Dephasing occurs when the difference between the phase velocity of the wake and the high-frequency photon pulse results in it slipping out of the accelerating refractive index gradient. A recent scheme for overcoming this restriction using an ionization front is dephasingless frequency shift using a ‘flying focus,’ a combination of a chirped laser pulse and an achromatic lens for spatiotemporal shaping of a laser pulse [21]. The flying focus was also used to mitigate the analogous process of electron beam dephasing in a plasma wakefield [89], in addition to related spatiotemporally structured focusing schemes [90, 91].

Another method for mitigating dephasing, in the context of electron acceleration is the use of tapered plasma density ramps [92, 93, 94, 95]. By having a non-uniform density, the plasma wavelength varies along the propagation length, which allows for locking the accelerating phase with the particle beam. Tapered density ramps were previously suggested as a way of increasing the frequency shifts in photon acceleration [8, 17].

This chapter develops an analytic model for dephasingless photon acceleration in the nonlinear plasma wake regime, based on a tapered plasma density profile. The model shows that if a wake can be sustained for arbitrary distance, there is no limit to the frequency shift achievable. The 1D model is used to design the density profile for quasi-3D particle-in-cell simulations for a broad driver beam, which demonstrates a frequency upshift of an 800 nm witness laser pulse to 80 nm by maintaining the phase matching between the wake and pulse centroid. The pulse experiences net energy gain of $5\times$ to 250 mJ and is compressed, leading to an ultrashort pulse of XUV radiation with relativistic intensity.

The rest of this chapter is outlined as follows. Section 3.1 presents background material on

photon acceleration and frequency shift. Section 3.2 presents the phase matching condition on which PA^∞ is based. Section 3.3 derives expressions for key locations and gradients in the wake behind the driver, dependent on quantities determined within the driver. Section 3.4 derives expressions for plasma fluid quantities at the end of the drive beam as functions of the drive beam parameters. Section 3.5 unites the work of the previous sections and completes the map from drive beam and unperturbed plasma density to plasma wake necessary for phase matched photon acceleration. The map is summarized and numerical verification is presented. Section 3.6 expands the PA^∞ model in the ultrashort and weak driver limits to find analytical expressions for the frequency shift and phase matching conditions. Using formulas obtained in section 3.6 or solving directly the differential system consisting of the frequency shift and phase matching conditions leads to the phase matched plasma density profile and expected frequency shift. This section mathematically demonstrates that the frequency shift is constrained only by the ability to maintain the wake, i.e. the photon acceleration is unlimited. Section 3.8 uses the PA^∞ tapered density profile in particle-in-cell (PIC) simulations to investigate the predicted frequency shift. Section 3.9 summarizes the results of the simulations and looks ahead to future research into photon acceleration.

3.1 Background : photon acceleration

A laser pulse that experiences a co-moving plasma density gradient will be upshifted in frequency. From eikonal solutions to the wave equation, well known ray-tracing solutions can be used for the temporal variation in the light [96, 16] to relate it to the density gradients generated in a wakefield. For a given dispersion relation D , the frequency ω and wavenumber k_z of an optical mode propagating in the z direction satisfy the following relations:

$$\begin{aligned}
dz/d\tau &= \partial D/\partial k_z, & dk_z/d\tau &= -\partial D/\partial z \\
dt/d\tau &= \partial D/\partial \omega, & d\omega/d\tau &= \partial D/\partial t,
\end{aligned}
\tag{3.1}$$

where τ parametrizes the ray path of the optical mode.

The optical mode will vary with distance propagated as

$$\frac{d\omega}{dz} = \frac{d\omega/d\tau}{dz/d\tau} = \frac{\partial D/\partial t}{\partial D/\partial k_z}.
\tag{3.2}$$

We assume that the laser pulse is propagating in a wakefield generated by a relativistic driver (either a relativistic particle beam or second laser pulse) propagating at velocity $v_d(z)$ and therefore change coordinates from (x, y, z, t) to $(x, y, z, \zeta = t - \int_0^z dz'/v_d(z'))$ [93]. The change in the frequency can, therefore, be expressed as

$$\frac{d\omega}{dz} = \frac{\partial D/\partial \zeta}{\partial D/\partial k_z}.
\tag{3.3}$$

For example, for linear plasma dispersion, $D = \omega^2 - \omega_p^2/\gamma - k_z^2 c^2$, and assuming that ions are immobile and the variations in the plasma density with respect to ζ are much larger than the variations in ω and k ,

$$\frac{d\omega}{dz} \simeq \frac{1}{2k_z c^2} \frac{\partial(\omega_p^2/\gamma)}{\partial \zeta}.
\tag{3.4}$$

Note that the γ factor in the linear dispersion relation is to allow for relativistically streaming electrons rather than, e.g. oscillations in the laser field, and therefore this dispersion relation is considered exact for a weak laser pulse.

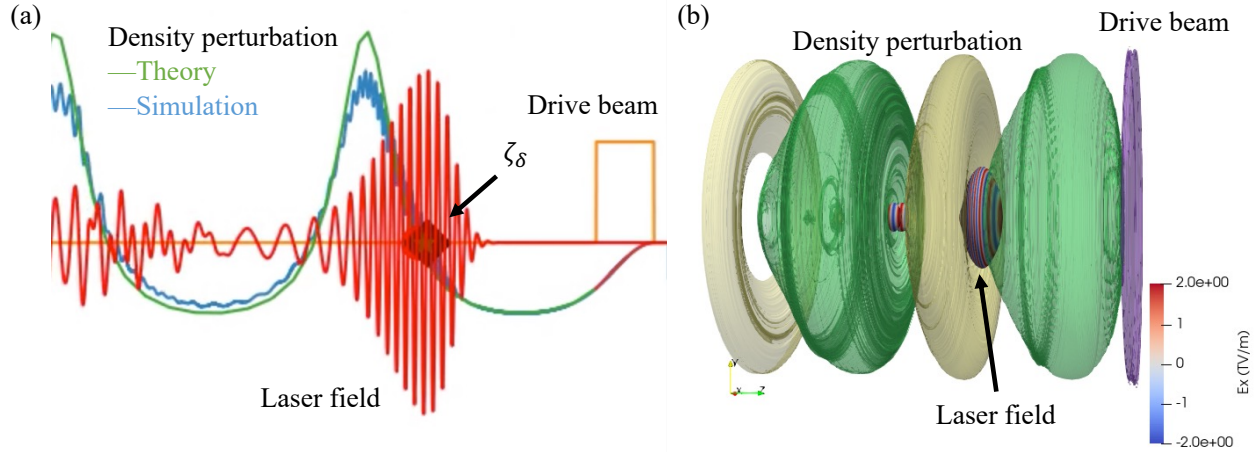


Figure 3.1: Schematic of unlimited photon acceleration. (a) Line drawing of PA^∞ layout, drive beam (in orange) propagates to the right through the plasma, laser pulse (red) co-propagates at the point ζ_δ where density (theory in green, simulation in blue) is increasing with distance ζ from the leading edge of the drive beam. (b) 3D visualization of PA^∞ , drive beam in purple propagates to the right through the plasma, laser pulse in blue and red co-propagates, wake indicated with isocontours of plasma density depression (in green) and elevation (in yellow).

3.2 Matching condition to mitigate dephasing

For positive frequency shift, the laser pulse must be at a phase in the wake where the density gradient is positive. However, the laser centroid moves at the group velocity of the laser pulse and so as the laser pulse shifts in frequency, its group velocity increases and the pulse will change position in the wake. To mitigate dephasing of the photon pulse, we use an ultrarelativistic electron beam to create the wake and a tapered density profile (similar to that proposed for mitigating dephasing in electron acceleration [93, 95, 97, 98]) to continuously increase the plasma wavelength and keep the laser pulse experiencing a positive plasma density gradient. The layout is shown in figure 3.1. A drive beam propagates to the right in the figures. Behind the drive beam, situated at the point in the wake where the plasma density is increasing with distance behind the drive beam, is the laser pulse. Not shown is the plasma density profile $n(z)$; obtaining n is a key part of this chapter.

For convenience, we choose the point where the density perturbation in the wake $\delta n =$

0 within the region of positive density gradient, hereby labelled as ζ_δ , as the location of the reference density gradient we are trying to track. This is not the maximum density gradient except in the linear regime, but the maximum is in general slightly behind ζ_δ , and using ζ_δ greatly simplifies the analysis. It can be shown that the maximum refractive index gradient occurs where the electric field of the wake is zero, which is close to where the density perturbation goes to zero. It can also be shown that the refractive index gradient is equal to the density gradient at the point where $\delta n = 0$ for a 1D wake.

To keep the pulse experiencing the greatest possible frequency shift, we require $\zeta_{\text{centroid}} = \zeta_\delta$ for all times, where ζ_{centroid} denotes the center of the witness laser pulse. Expressed in differential form, we have

$$\frac{d\zeta_{\text{centroid}}}{dz} = \frac{d\zeta_\delta}{dz}, \quad (3.5)$$

which may be written as

$$\frac{1}{v_{\text{centroid}}} - \frac{1}{v_d} = \frac{d\zeta_\delta}{dn} \frac{dn}{dz}. \quad (3.6)$$

For an ultrarelativistic particle beam driver, we make the approximation $v_d \rightarrow c$. Assuming that the laser pulse moves at the linear group velocity and the plasma is underdense, $\omega_p^2/\omega^2 \ll 1$, we obtain an equation relating the z profile of the plasma number density to the variation in the wake position of the zero density perturbation ζ_δ with plasma density,

$$\frac{dn}{dz} \simeq \frac{1}{2c} \frac{\omega_p^2}{\omega^2} \left[\frac{d\zeta_\delta}{dn} \right]^{-1}. \quad (3.7)$$

Note that $\gamma(\zeta_\delta) = 1$, where the density perturbation is zero.

The phase matching condition in equation (3.7) and the frequency shift relation in equation (3.4) together are a coupled pair of differential equations for the plasma density profile that will give phase matched photon acceleration as well as the predicted frequency shift given the phase matched density profile. To close them, it remains to determine the depen-

dence of the quantities $\left. \frac{\partial(\omega_p^2/\gamma)}{\partial\zeta} \right|_{\zeta_\delta}$ and $\frac{d\zeta_\delta}{dn}$ on n , ω , and possibly z .

3.3 Description of the wake behind the driver

In this section, the position ζ_δ in the wake where there is a positive density gradient and $\delta n = 0$, and the density gradient there, $\left. \frac{\partial(\omega_p^2/\gamma)}{\partial\zeta} \right|_{\zeta_\delta}$, are obtained. The wake behind the driver, the region indicated by the dashed black box in Fig. 3.2, is modeled with the Akhiezer-




Figure 3.2: PA^∞ requires analytic expressions for position ζ_δ of zero perturbation and increasing plasma density and for the gradient there, $\left. \frac{\partial(\omega_p^2/\gamma)}{\partial\zeta} \right|_{\zeta_\delta}$. These expressions are obtained by solving the plasma fluid equations behind the driver, the region indicated by the dashed black box here. The driver density is the orange line, creating the plasma response shown in blue. The analytic expression for the plasma response in the absence of the witness pulse is in green. Information about the fluid quantities at the end of the driver, where $\zeta = \zeta_d$, is also needed, but this is determined in section 3.4.

Polovin (AP) [99] relativistic wave. This is a nonlinear plasma wave with a given amplitude γ_m and velocity $\beta = v_p/c$. The Lagrangian form of the plasma equations describing the AP wave is used here to develop a solution parametrized in terms of a periodic coordinate φ as derived by Infeld and Rowland in [100] (particularly, the expression of the Infeld-Rowland-Akhiezer-Polovin wave derived in Verma et al. [101, 102].) With the parametrization of the wake behind the driver, the position of zero density perturbation, ζ_δ , and the plasma density gradient there, $\left. \frac{\partial(\omega_p^2/\gamma)}{\partial\zeta} \right|_{\zeta_\delta}$, are found as functions of the wake parameters φ_{st} and amplitude

γ_m defining the wake behind the driver. Then, the dependence of the starting phase φ_{st} on the wake within the driver is found, particularly on the fluid Lorentz factor γ_d , momentum p_d , and electric field E_{zd} at the end of the driver. The dependence of the wake amplitude γ_m on fluid quantities in the wake will be determined in the subsequent section, section 3.4.

Under the assumption that $d \log(n)/dz \ll k_p(z)$ for all z , i.e. the plasma density gradients are long compared to the wake period scale, to lowest order the wake may be assumed to follow the uniform plasma solution with local density $n(z)$. Hence, the density in the wake $n_w(\zeta; z)$ is a function of ζ and is parametrized by z , i.e. by the local density $n(z)$.

In a 1D model the density in the wake $n_w(\zeta; z)$, together with the fluid momentum p and velocity v and electric field E_z , satisfies the cold unmagnetized fluid equations

$$p = \gamma m v, \quad \gamma = (1 - v^2/c^2)^{-1/2}, \quad (3.8a)$$

$$\frac{\partial n_w}{\partial t} + \frac{\partial n_w v}{\partial z} = 0, \quad (3.8b)$$

$$\frac{\partial p}{\partial t} + v \frac{\partial p}{\partial z} = q E_z, \quad (3.8c)$$

$$\frac{\partial E_z}{\partial z} = 4\pi q (n_w - n), \quad (3.8d)$$

$$0 = \frac{4\pi}{c} q n_w v + \frac{1}{c} \frac{\partial E_z}{\partial t}. \quad (3.8e)$$

The plasma particles are electrons with charge $q = -e$ and mass $m = m_e$ and there is a static neutralizing background of positive ions. The last two equations, Gauss' law (3.8d) and Ampère's law (3.8e), combine to give

$$\frac{\partial E_z}{\partial t} + v \frac{\partial E_z}{\partial z} = -4\pi q n v. \quad (3.9)$$

Performing a Lagrangian change of coordinates,

$$z(\alpha, \tau) = z_0(\alpha) + \int_0^\tau v(\alpha, \tau') d\tau', \quad t(\alpha, \tau) = \tau, \quad (3.10)$$

leads to the Lagrangian form of the fluid equations,

$$n_w(\alpha, \tau) = \frac{n_w(\alpha, 0) \frac{\partial z_0}{\partial \alpha}(\alpha)}{\frac{\partial z_0}{\partial \alpha}(\alpha) + \int_0^\tau \frac{\partial v}{\partial \alpha}(\alpha, \tau') d\tau'} , \quad (3.11a)$$

$$\frac{\partial p}{\partial \tau} = qE_z , \quad (3.11b)$$

$$\frac{\partial E_z}{\partial \tau} = -4\pi q n_0 \frac{p}{m\gamma} . \quad (3.11c)$$

The last two equations combine to form the equation for the relativistic harmonic oscillator

$$\frac{\partial^2 p}{\partial \tau^2} = -\frac{4\pi q^2 n p}{m\gamma} = -\omega_p^2 \frac{p}{\gamma} , \quad (3.12)$$

where $\omega_p^2 = \frac{4\pi q^2 n}{m}$.

The relativistic harmonic oscillator equation can be directly integrated:

$$\begin{aligned} \int_0^\tau \frac{\partial^2 p}{\partial \tau'^2} \frac{\partial p}{\partial \tau'} d\tau' &= -\omega_p^2 \int_0^\tau \frac{p}{\sqrt{1 + p^2/m^2 c^2}} \frac{\partial p}{\partial \tau'} d\tau' \\ &= 2m^2 c^2 \omega_p^2 (B - \gamma) , \end{aligned} \quad (3.13)$$

where B is a constant of integration. In the last line, it must be that $B \geq \gamma(\alpha, \tau)$, which implies that $B(\alpha) = \gamma_m(\alpha) = \sqrt{1 + p_m(\alpha)^2/m^2 c^2}$. While the constant of integration can be distinct for each α , the solutions for a traveling wave will see a uniform γ_m so the α can be suppressed.

The units are removed with the substitutions $p \rightarrow p/mc, \tau \rightarrow \omega_p \tau$, to obtain

$$\frac{\partial p}{\partial \tau} = \pm \sqrt{2} \sqrt{\gamma_m - \sqrt{1 + p^2}} . \quad (3.14)$$

This integral separates for direct integration:

$$\int \pm \sqrt{2} d\tau = \int \frac{dp}{\sqrt{\gamma_m - \sqrt{1 + p^2}}} . \quad (3.15)$$

At this point the change of variables

$$(\gamma_m - 1) \sin^2 \varphi = \gamma_m - \sqrt{1 + p^2}, \quad (\gamma_m - 1) \sin 2\varphi d\varphi = -\frac{1}{\sqrt{1 + p^2}} p dp, \quad (3.16)$$

leads to the solution [100, 101]

$$\tau = \sqrt{2(\gamma_m + 1)} E(\varphi, \kappa^2) - \sqrt{\frac{2}{\gamma_m + 1}} F(\varphi, \kappa^2) + D, \quad (3.17)$$

where F, E are the incomplete elliptic integrals of the first and second kind (note E is not the electric field), $\kappa^2 = \frac{\gamma_m - 1}{\gamma_m + 1}$, $\kappa'^2 = \frac{2}{\gamma_m + 1}$, and $D(\alpha)$ is another constant of integration.

Verma et. al. [101] showed that a traveling wave is obtained if γ_m is independent of α and if $D(\alpha) = \omega_p \alpha / v_p$ and $v_p = c\beta$ is the wave phase velocity. The AP wave can be expressed in Lagrangian coordinates α, τ as:

$$z(\alpha, \tau) = z_0(\alpha) + \Delta z(\varphi(\alpha, \tau)), \quad (3.18a)$$

$$\Delta z(\varphi) = \frac{c}{\omega_p} \frac{2\kappa}{\kappa'} \sin(\varphi), \quad (3.18b)$$

$$\gamma(\varphi) = \gamma_m - (\gamma_m - 1) \sin^2(\varphi), \quad (3.18c)$$

$$p(\varphi) = \cos \varphi \sqrt{\gamma_m - 1} \sqrt{1 + \gamma}, \quad (3.18d)$$

$$qE_z(\varphi) = \frac{\partial p}{\partial \tau} = \sin \varphi \sqrt{2(\gamma_m - 1)}, \quad (3.18e)$$

$$\omega_p \tau = \frac{2}{\kappa'} E(\varphi, \kappa^2) - \kappa' F(\varphi, \kappa^2) + \frac{\omega_p \alpha}{v_p}. \quad (3.18f)$$

The last equation determines φ implicitly as a function of α, τ .

The AP solution to the cold relativistic fluid equations is a traveling wave that can be expressed in the single parameter $\zeta = t - z/v_p$. To obtain the traveling wave solution, the condition $z_0(\alpha) = \alpha$ is imposed. Expressed in Eulerian coordinates, the relativistic traveling wave has the form

$$\omega_p \zeta = \frac{2}{\kappa'} E(\varphi, \kappa^2) - \kappa' F(\varphi, \kappa^2) - \frac{1}{\beta} \frac{2\kappa}{\kappa'} \sin \varphi \quad (3.19a)$$

$$\gamma(\varphi) = \gamma_m - (\gamma_m - 1) \sin^2(\varphi) , \quad (3.19b)$$

$$p(\gamma) = \cos \varphi \sqrt{\gamma_m - 1} \sqrt{1 + \gamma} , \quad (3.19c)$$

$$qE_z(\varphi) = \frac{\partial p}{\partial \tau} = \sin \varphi \sqrt{2(\gamma_m - 1)} , \quad (3.19d)$$

$$n_w(\gamma) = \frac{n}{1 - p(\gamma)/\gamma} , \quad (3.19e)$$

where parameters $\varphi = 2j\pi$, $j = \dots, -2, -1, 0, 1, 2, \dots$ correspond to maxima in the plasma density and $\varphi = j + \pi/2$ correspond to points of zero perturbation, i.e. $n_w = n$. Furthermore, we demonstrate that $\varphi = (2j + 3/2)\pi$ corresponds to points of zero perturbation and increasing plasma density, $\partial n_w / \partial \zeta > 0$. These equations completely describe the wake behind the drive electron beam, given the mapping from drive beam to amplitude γ_m and starting parameter φ_{st} .

With the wake behind the driver completely described in terms of amplitude γ_m and parameter φ_{st} , the quantities ζ_δ and $\left. \frac{\partial(\omega_p^2/\gamma)}{\partial \zeta} \right|_{\zeta_\delta}$ can be derived. The position ζ_δ , measured from the leading edge of the drive beam, may be expressed as

$$\zeta_\delta = \zeta(3\pi/2) - \zeta(\varphi_{st}) + L_d/c , \quad (3.20)$$

where $\zeta(\varphi)$ is the position in the undriven wake as a function of φ , L_d is the length of the drive beam, and φ_{st} is determined so that the wake is continuous across the transition from drive beam to undriven wake. To enforce continuity of the wake, we will require that electric field and momentum are continuous at the transition from within the drive beam to behind the drive beam. Let γ_d denote the plasma Lorentz factor at the end of the drive beam, p_d be the plasma momentum at the end of the drive beam, and E_{zd} the electric field at the end of the driver. The equations for p and E_z behind the driver can be arranged to determine φ

in terms of fluid quantities:

$$\cos \varphi = \frac{p}{\sqrt{\gamma_m - 1} \sqrt{1 + \gamma}}, \quad (3.21a)$$

$$\sin \varphi = \frac{qE_z}{\sqrt{2(\gamma_m - 1)}}. \quad (3.21b)$$

These equations combine to determine $\tan \varphi_{st}$ as a function of the fluid quantities:

$$\begin{aligned} \tan \varphi_{st} &= \frac{\sin \varphi_{st}}{\cos \varphi_{st}} = \frac{qE_{zd}}{\sqrt{2(\gamma_m - 1)}} \cdot \frac{\sqrt{\gamma_m - 1} \sqrt{1 + \gamma_d}}{p_d} \\ &= \frac{qE_{zd}}{p_d} \sqrt{\frac{1 + \gamma_d}{2}}. \end{aligned} \quad (3.22)$$

Now φ_{st} is determined by the fluid quantities within the drive beam as well as γ_m . The dependence of these on the drive beam parameters will be derived in the next section.

Next, the dependence of $\left. \frac{\partial(\omega_p^2/\gamma)}{\partial\zeta} \right|_{\zeta_\delta}$ on the unperturbed plasma density n and the wake amplitude γ_m is determined. Because ω_p^2 is proportional to n_w with constant of proportionality $e^2/m_e\varepsilon_0$, it suffices to calculate $\partial(n_w/\gamma)/\partial\zeta|_{\zeta_\delta}$. Because evolution in z is considered to be much slower than evolution in ζ , in this section n , ω_p , $\gamma_m(z)$, β , κ , and κ' are treated as constants. Hence, all quantities ultimately depend only on ζ and we will derive $d(n_w/\gamma)/d\zeta|_{\zeta_\delta}$. When $\delta n = 0$, then $n_w = n$ so $\beta = p = 0$, $\gamma = 1$, and $\varphi = (j + 1/2)\pi$, $j = \dots, -2, -1, 0, 1, \dots$

First, we'll show that $\left. \frac{d(n_w/\gamma)}{d\zeta} \right|_{\zeta_\delta} = \left. \frac{dn_w}{d\zeta} \right|_{\zeta_\delta}$, by observing that $\left. \frac{d\gamma}{d\zeta} \right|_{\zeta_\delta} = 0$:

$$\begin{aligned} \left. \frac{d\gamma}{d\zeta} \right|_{\zeta_\delta} &= \left. \frac{d\gamma}{d\varphi} \right|_{\varphi=(n+1/2)\pi} \left. \frac{d\varphi}{d\zeta} \right|_{\zeta_\delta} \\ &= -(\gamma_m - 1) \sin([2j + 1]\pi) \frac{1}{\omega_p \left(\frac{2}{\kappa'} \sqrt{1 - \kappa^2 \sin^2 \varphi} - \frac{\kappa'}{1 - \kappa^2 \sin^2(2n + \pi)} - 2 \frac{\kappa}{\kappa'} \sin^2(2n + 1)\pi \right)} \\ &= -(\gamma_m - 1) \sin([2j + 1]\pi) \frac{\omega_p \kappa'}{\kappa' + 2\kappa} = 0. \end{aligned} \quad (3.23)$$

Then, because $\gamma_m(\varphi = (j + 1/2)\pi) = 1$, we have

$$\left. \frac{d(n_w/\gamma)}{d\zeta} \right|_{\zeta_\delta} = \left. \frac{dn_w}{d\zeta} \right|_{\zeta_\delta}. \quad (3.24)$$

We evaluate $\left. \frac{dn_w}{d\zeta} \right|_{\zeta_\delta}$ by the chain rule,

$$\left. \frac{dn_w}{d\zeta} \right|_{\zeta_\delta} = \left. \frac{dn}{d\beta} \right|_{\beta=0} \left. \frac{d\beta}{d\varphi} \right|_{\varphi=(j+1/2)\pi} \left. \frac{d\varphi}{d\zeta} \right|_{\zeta_\delta}. \quad (3.25)$$

Writing the wake plasma density in terms of $\beta = p/\gamma$, $n_w = n(z)/(1 - \beta)$, we calculate $dn_w/d\beta$ and evaluate at $\beta = 0$:

$$\begin{aligned} \frac{dn_w}{d\beta} &= \frac{n(z)}{(1 - \beta)^2}, \\ \left. \frac{dn_w}{d\beta} \right|_{\beta=0} &= n(z). \end{aligned} \quad (3.26)$$

Then we write the wave velocity in terms of the Lorentz factor γ and use the relation for $\gamma(\varphi)$:

$$\begin{aligned} \frac{d\beta}{d\varphi} &= \left(\frac{\partial\beta}{\partial p} \frac{dp}{d\gamma} + \frac{\partial\beta}{\partial\gamma} \right) \frac{d\gamma}{d\varphi} \\ &= -\frac{2 \sin \varphi \sqrt{\gamma_m - 1}}{\sqrt{\gamma + 1}}, \\ \left. \frac{d\beta}{d\varphi} \right|_{\varphi=(j+1/2)\pi} &= -\sin \varphi \sqrt{2(\gamma_m - 1)}. \end{aligned} \quad (3.27)$$

Since $\omega_p \zeta = \frac{2}{\kappa'} E(\varphi, \kappa^2) - \kappa' K(\varphi, \kappa^2) - \frac{2\kappa}{\beta\kappa'} \sin \varphi$, we have

$$\begin{aligned} \omega_p \frac{d\zeta}{d\varphi} &= \frac{2}{\kappa'} \sqrt{1 - \kappa^2 \sin^2 \varphi} - \kappa' \frac{1}{1 - \kappa^2 \sin^2 \varphi} - \frac{2\kappa}{\beta\kappa'} \cos \varphi \\ &= \frac{2}{\kappa'} \sqrt{1 - \kappa^2} - \kappa' \frac{1}{1 - \kappa^2}, \\ \omega_p \frac{d\zeta}{d\varphi} \Big|_{\varphi=(j+1/2)\pi} &= \frac{2}{\kappa'} \kappa' - \kappa' \frac{1}{\kappa'} = 1. \end{aligned} \quad (3.28)$$

Altogether, the plasma gradient where $\delta n = 0$ is

$$\left. \frac{dn_w}{d\zeta} \right|_{\zeta_\delta} = n(z)(-\sin(\varphi))\sqrt{2(\gamma_m - 1)}\omega_p(z), \quad \varphi = (j + 1/2)\pi. \quad (3.29)$$

The positive plasma gradient required for positive frequency shift then means that ζ_δ occurs at $\varphi = (2j + \frac{3}{2})\pi$. We can therefore write

$$\left. \frac{d(\omega_p^2/\gamma)}{d\zeta} \right|_{\zeta_\delta} = \omega_p(z)^3 \sqrt{2(\gamma_m - 1)}. \quad (3.30)$$

In this section, we have modeled the wake behind the driver as a traveling relativistic plasma wave. We presented a periodic parametrization of this wave and derived from this representation quantities ζ_δ and $\left. \frac{\partial \omega_p^2/\gamma}{\partial \zeta} \right|_{\zeta_\delta}$ necessary for understanding frequency shift and phase matching. It remains to determine the dependence of the amplitude γ_m on the drive beam, as well as the dependence of the electric field E_{zd} , fluid momentum p , and fluid gamma factor γ_d at the end of the drive beam on the drive beam parameters.

3.4 Description of the wake within the driver

In this section we determine the map from drive beam parameters to wake behind the driver. In the previous section we found that the wake behind the driver can be completely understood in terms of its amplitude γ_m and a continuity condition between the wake within the drive beam and the wake behind the drive beam. The continuity condition depends on the plasma particle Lorentz factor γ_d , momentum p_d , and electric field E_{zd} at the end of the drive beam. In this section we derive γ_m , γ_d , p_d , and E_{zd} in terms of the drive beam parameters. To determine the wake quantities at the end of the drive beam, we solve the fluid equations within the drive beam, the region indicated the dashed black box in Fig. 3.3, where there is a nonzero drive beam density contribution to the plasma density.

This section consists of three parts. First, the plasma fluid equations within the drive beam are presented. Next, these are solved in three cases depending on the relation between the drive beam density n_d and the unperturbed plasma density $n = n(z)$, namely when $n_d < n/2$, $n_d = n/2$, and $n_d > n/2$. Since there are three solution sets, the quantities γ_d , p_d ,




Figure 3.3: PA^∞ requires analytic expressions for plasma fluid quantities at the end of the driver, $\zeta = \zeta_d$. These expressions are obtained by solving the plasma fluid equations within the driver, the region indicated by the dashed black box in this figure. The driver density is the orange line, creating the plasma response within the driver shown in red (not to be confused with the oscillatory laser pulse further back in the wake, also in red).

and E_{zd} are expressed in three different cases. The electric potential at the end of the driver, ϕ_d , will also be presented in each case, since it is needed for determining the amplitude γ_m of the wake behind the driver. This section ends by deriving the relation $\gamma_m = 1 - n_d\phi_d$ between the drive beam parameters and the amplitude of the wake behind the driver.

We begin with the plasma fluid equations within the drive beam. We assume an ultra-relativistic ($\beta_d \rightarrow 1$) drive beam, in which case the fluid equations are

$$\frac{n_w}{n} = \frac{1}{1 - \beta}, \quad (3.31a)$$

$$\frac{1}{\omega_p^2} \frac{\partial^2}{\partial \zeta^2} (\gamma(1 - \beta)) = \frac{\beta}{1 - \beta} + n_d/n. \quad (3.31b)$$

Using the approximation that the plasma density $n(z)$ is slowly varying, we may use the wake solutions in constant density to find $n_w(\zeta; n(z))$. We therefore treat $n(z)$ as a constant, n , and write $\zeta = \omega_p \zeta$, $n_d = n_d/n$, and $L_d = \omega_p L_d/c$.

Following Rosenzweig [103], we let

$$x(\zeta) \equiv \gamma(1 - \beta) = \sqrt{\frac{1 - \beta}{1 + \beta}}. \quad (3.32)$$

This leads to

$$\frac{\partial^2 x}{\partial \zeta^2} = \frac{1}{2} \left(\frac{1}{x^2} - 1 + 2n_d \right) . \quad (3.33)$$

We assume that the plasma is quiescent ahead of the driver, so for $\zeta \leq 0$, $\beta = 0 = p = E$, and $n = 1 = \gamma = x$.

This integrates once to give

$$\left(\frac{\partial x}{\partial \zeta} \right)^2 = 2(1 - n_d) - 1/x + (2n_d - 1)x . \quad (3.34)$$

We make a detour to discuss the electric field and the choice of sign for the square root that equals $\frac{\partial x}{\partial \zeta}$. Recall that Gauss' law in this ultrarelativistic problem states

$$\begin{aligned} -\frac{1}{v_d} \frac{\partial E_z}{\partial \zeta} &= 4\pi e(n_i - n_e - n_d) \\ \frac{\partial E_z}{\partial \zeta} &= -(1 - n_w - n_d) = \frac{\partial^2 x}{\partial \zeta^2} . \end{aligned} \quad (3.35)$$

We can make the identifications $E_z = \pm \frac{\partial x}{\partial \zeta}$ and $\phi = \pm(-x + 1)$ since $E_z = 0 = \frac{\partial x}{\partial \zeta}$ for $\zeta \leq 0$. Since we made the assumption that the plasma species and beam particles are electrons, immediately behind the driver the electron density is decreased; then $\frac{\partial E_z}{\partial \zeta} = n_w + n_d - 1 > 0$ and we start with the positive square root for $\frac{\partial x}{\partial \zeta}$.

For identifying dependence of fluid quantities on drive beam parameters we first determine the dependence on x :

$$\gamma = \frac{1 + x^2}{2x} , \quad (3.36a)$$

$$p = \frac{1 - x^2}{2x} , \quad (3.36b)$$

$$E_z = \frac{\partial x}{\partial \zeta} = \text{sgn}(E_z) \sqrt{2(1 - n_d) - 1/x + (2n_d - 1)x} , \quad (3.36c)$$

$$\phi = 1 - x . \quad (3.36d)$$

The term $(2n_d - 1)x$ changes sign based on the relation of n_d and $1/2$, as does the qualitative behavior of the solution and the technique used to derive the solution. There are three cases

to consider:

1. $0 < n_d < 1/2$
2. $n_d = 1/2$
3. $n_d > 1/2$

Case 1: $0 < n_d < 1/2$

This is the case studied in Rosenzweig [103] but we include the derivation for completeness. We will find that x is oscillatory in ζ where as in the other cases x is monotone.

The equation for x is

$$\frac{\partial x}{\partial \zeta} = \sqrt{1 + 1 - 2n_d - 1/x - (1 - 2n_d)x} . \quad (3.37)$$

Following Rosenzweig, we rearrange the equation and integrate:

$$\sqrt{1 - 2n_d} \int_1^x \frac{\partial \zeta}{\partial x'} dx' = \int_1^x \sqrt{\frac{x'}{(x' - 1)(1/(1 - 2n_d) - x')}} dx' . \quad (3.38)$$

Let

$$(x' - 1)(1 - 2n_d) \equiv t^2, \quad (3.39)$$

so $dx'(1 - 2n_d) = 2tdt'$. We have

$$\begin{aligned} \sqrt{1 - 2n_d}(\zeta(x)) &= \int_0^{t(x)} \sqrt{\frac{1 + t'^2/(1 - 2n_d)}{\frac{t'^2}{1 - 2n_d} \left(\frac{1}{1 - 2n_d} - 1 + 1 - x' \right)}} \frac{2t' dt'}{1 - 2n_d} \\ &= \frac{2}{\sqrt{1 - 2n_d}} \int_0^{t(x)} \sqrt{\frac{1 - 2n_d + t'^2}{2n_d - t'^2}} dt' . \end{aligned} \quad (3.40)$$

Now we let

$$\frac{t'}{\sqrt{2n_d}} = \cos \theta , \quad (3.41a)$$

$$dt' = -\sqrt{2n_d} \sin \theta d\theta , \quad (3.41b)$$

and arrive at

$$\begin{aligned}
(1 - 2n_d)\zeta(x) &= 2 \int_0^{t(x)} \sqrt{\frac{1 - 2n_d + 2n_d \cos^2 \theta}{2n_d - 2n_d \cos^2 \theta}} (-\sqrt{2n_d} \sin \theta d\theta) \\
&= 2 \int_\theta^{\pi/2} \sqrt{1 - 2n_d \sin^2 \theta} d\theta .
\end{aligned} \tag{3.42}$$

We can write this expression in terms of the incomplete elliptic integral of the second kind,

$$E(k, \varphi) \equiv \int_0^\varphi \sqrt{1 - k^2 \sin^2 \theta} d\theta , \tag{3.43}$$

and the complete elliptic integral of the second kind,

$$E(k) = E(k, \pi/2). \tag{3.44}$$

This gives an implicit relation for θ as a function of ζ :

$$(1 - 2n_d)\zeta(\theta) = 2(E(\sqrt{2n_d}) - E(\sqrt{2n_d}, \theta)). \tag{3.45}$$

Recall that the above relation is valid within the driver, for $-L_d \leq \zeta \leq 0$. Solving for $\theta(\zeta)$, we can find x :

$$x(\zeta) = 1 + \frac{2n_d}{1 - 2n_d} \cos^2 \theta(\zeta) . \tag{3.46}$$

Here $x(\zeta)$ is oscillatory, so determining the electric field requires determining sign changes.

Since the electric field is equal to $\pm \frac{\partial x}{\partial \zeta}$, $\frac{\partial x}{\partial \theta} = \frac{2n_d}{1 - 2n_d} \sin 2\theta(\zeta)$, and

$$\frac{\partial \theta}{\partial \zeta} = \frac{2n_d - 1}{2\sqrt{1 - n_d \sin^2 \theta}} > 0 , \tag{3.47}$$

we know that

$$\text{sgn}(E_z) = \text{sgn}(\sin(2\theta)) . \tag{3.48}$$

In summary, x_d is a function of an intermediate variable θ_d which is determined implicitly

from the following equations:

$$(1 - 2n_d)L_d = 2 \left(E(\sqrt{2n_d}) - E(\theta_d, \sqrt{2n_d}) \right) , \quad (3.49a)$$

$$x_d = 1 + \frac{2n_d}{1 - 2n_d} \cos^2 \theta_d , \quad (3.49b)$$

$$\text{sgn}(E_{zd}) = \text{sgn}(\sin(2\theta_d)) . \quad (3.49c)$$

Case 2: $n_d = 1/2$

In this case, the equation for x simplifies to

$$\frac{\partial x}{\partial \zeta} = \sqrt{1 - \frac{1}{x}} . \quad (3.50)$$

Since $\frac{\partial \zeta}{\partial x} > 0$ when $x = 1$, and $\zeta = 0$ and $\frac{\partial x}{\partial \zeta} > 0$ for all ζ , it must be that $x(\zeta)$ is a positive, monotonically increasing function.

Integrating, we have

$$\zeta = \int_1^x \sqrt{\frac{x'}{x' - 1}} dx' . \quad (3.51)$$

Performing a change of variables

$$x - 1 = \sinh^2 \theta, \quad dx = 2 \sinh(\theta) \cosh(\theta) d\theta , \quad (3.52)$$

leads to

$$\begin{aligned} \zeta &= \int_0^\theta 2\sqrt{1 + \sinh^2 \theta} \cosh \theta d\theta \\ &= \sinh^{-1} \sqrt{x - 1} + \sqrt{x - 1} \sqrt{x} . \end{aligned} \quad (3.53)$$

Since $x(\zeta)$ is non-oscillatory, E_z and n_w will be monotonic and in particular, E_z stays positive within the drive beam. Similarly, the electric potential ϕ is negative within the drive beam.

The quantity x_d is determined implicitly from the equation

$$L_d = \sinh^{-1} \sqrt{x_d - 1} + \sqrt{x_d - 1} \sqrt{x_d} \quad (3.54)$$

and $\text{sgn}(E_{zd}) = +1$.

Case 3: $n_d > 1/2$

This case is not considered in prior works, that we are aware of. The equation for x is

$$\frac{\partial x}{\partial \zeta} = \sqrt{2(1 - n_d) - 1/x + (2n_d - 1)x} . \quad (3.55)$$

Note that this is positive and increasing with x , so x is a positive, monotonically increasing function of ζ . Rearranging so we can solve for $\zeta(x)$:

$$\frac{\partial \zeta}{\partial x} = \frac{1}{\sqrt{[(2 - 2n_d)x - 1 + (2n_d - 1)x^2] / x}} \quad (3.56)$$

$$\sqrt{2n_d - 1} \frac{\partial \zeta}{\partial x} = \frac{\sqrt{x}}{\sqrt{(x - 1) \left(x + \frac{1}{2n_d - 1}\right)}} . \quad (3.57)$$

Integrating, we have

$$\sqrt{2n_d - 1} \int_0^\zeta d\zeta' = \int_1^x \frac{\sqrt{x'}}{\sqrt{(x' - 1) \left(x' + \frac{1}{2n_d - 1}\right)}} dx' . \quad (3.58)$$

Following a similar vein to Rosenzweig's solution technique for $n_d < 1/2n$ [103], we use the substitution:

$$(x - 1)(2n_d - 1) = t^2, \quad (2n_d - 1)dx = 2tdt , \quad (3.59)$$

after which the integral becomes

$$\sqrt{2n_d - 1}\zeta = \frac{2}{\sqrt{2n_d - 1}} \int_0^t \frac{\sqrt{t^2 + 2n_d - 1}}{\sqrt{t^2 + 2n_d}} dt . \quad (3.60)$$

This integral can be found in Byrd and Friedman [104][221.04,313.02]:

$$\int_0^t \frac{\sqrt{t'^2 + b^2}}{\sqrt{t'^2 + a^2}} dt' = \frac{gb^2}{k'^2} \left[k'^2 F \left(\tan^{-1} \left(\frac{t}{b} \right), k \right) - E \left(\tan^{-1} \left(\frac{t}{b} \right), k \right) + \sqrt{1 - k^2 \sin^2 \left(\tan^{-1} \left(\frac{t}{b} \right), k \right)} \frac{t}{b} \right], \quad (3.61)$$

where

$$\begin{aligned} k^2 &= \frac{a^2 - b^2}{a^2} = \frac{1}{2n_d}, & k'^2 &= 1 - k^2 = \frac{2n_d - 1}{2n_d}, \\ a^2 &= 2n_d, & b^2 &= 2n_d - 1, & g &= \frac{1}{a} = \frac{1}{\sqrt{2n_d}}. \end{aligned} \quad (3.62)$$

To compare the solution for $n_d > 1/2$ with the other cases, an angle substitution can be employed,

$$\tan \theta = \frac{t}{b}. \quad (3.63)$$

With this substitution, the integral becomes

$$(2n_d - 1)\zeta = 2 \left(gb^2 \left\{ \frac{1}{k'^2} \left[k'^2 F(\theta, k) - E(\theta, k) + \sqrt{1 - k^2 \sin^2 \theta} \tan \theta \right] \right\} \right). \quad (3.64)$$

Observing that

$$\frac{gb^2}{k'^2} = \sqrt{2n_d}, \quad (3.65)$$

the relation for ζ can be expressed as

$$(2n_d - 1)\zeta = 2\sqrt{2n_d} \left(k'^2 F(\theta, k) - E(\theta, k) + \tan \theta \sqrt{1 - k^2 \sin^2 \theta} \right). \quad (3.66)$$

In summary, x_d is a function of θ_d , which is implicitly a function of n_d and L_d , and $E_{zd} > 0$:

$$(2n_d - 1)L_d = 2\sqrt{2n_d} \left(k'^2 F(\theta_d, k) - E(\theta_d, k) + \tan \theta_d \sqrt{1 - k^2 \sin^2 \theta_d} \right), \quad (3.67a)$$

$$x_d = 1 + \tan^2 \theta_d, \quad (3.67b)$$

$$\text{sgn } E_{zd} = +1 . \quad (3.67c)$$

3.4.1 Discussion of the wake solutions within the driver

In each of these cases, the equations (3.36) can be used to determine the fluid Lorentz factor γ , the plasma electric field E_z , and the fluid momentum p as functions of x within the driver. The dependence of x on ζ is typically implicit. In each case, the dependence of x_d and the sign of E_{zd} on n_d and L_d is determined as detailed in equations (3.67), (3.54), and (3.49). With γ_d , p_d , and E_{zd} , continuity condition ϕ_{st} can be evaluated using equation (3.22) as determined in section 3.3.

Next the amplitude of the wake behind the driver γ_m is derived along with its dependence on drive beam parameters. This derivation was provided by Rosenzweig [103]. Behind the wake, $n_d = 0$, and equation (3.33) becomes

$$\frac{\partial^2 x}{\partial \tau^2} - \frac{1}{2} \left(\frac{1}{x^2} - 1 \right) = 0 . \quad (3.68)$$

Integrating this equation, we obtain

$$\frac{1}{2} \left(\frac{\partial x}{\partial \tau} \right)^2 + \frac{1}{2} \left(\frac{1}{x} + x \right) = C \quad (3.69)$$

for C a constant. Recall that $x = \sqrt{\frac{1-\beta}{1+\beta}}$ so

$$\frac{1}{x} + x = \sqrt{\frac{1+\beta}{1-\beta}} + \sqrt{\frac{1-\beta}{1+\beta}} \quad (3.70)$$

$$= 2\gamma \quad (3.71)$$

and

$$\frac{1}{2} \left(\frac{\partial x}{\partial \tau} \right)^2 + \gamma = C. \quad (3.72)$$

Since we know from section 3.3 that the wake behind the driver is oscillatory, we achieve

$\frac{\partial x}{\partial \tau} = 0$ and we can make the identification $C = \gamma_m$.

By the continuity of the wake,

$$\begin{aligned}
\gamma_m &= \frac{1}{2} \left(\left. \frac{\partial x}{\partial \zeta} \right|_{L_d} \right)^2 + \gamma_d \\
&= \left\{ 1 - n_d - \frac{1}{2x_d} + \left(n_d - \frac{1}{2} \right) x_d \right\} + \left\{ \frac{1}{2} \left(\frac{1}{x_d} + x_d \right) \right\} \\
&= 1 - n_d + n_d x_d .
\end{aligned} \tag{3.73}$$

If we recall that $\phi_d = 1 - x_d$, we can express the wake amplitude in terms of fluid quantities just as we did φ_{st} . Just as with the other fluid quantities γ_d , p_d , and E_{zd} , we know ϕ_d in terms of x_d and x_d implicitly in terms of drive beam parameters n_d , L_d . Hence, the equation for γ_m ,

$$\gamma_m = 1 - n_d \phi_d , \tag{3.74}$$

is the final equation needed to describe ζ_d and $\left. \frac{\partial \delta n}{\partial \zeta_\delta} \right|_{\zeta_\delta}$ in terms of the drive beam parameters.

Using the solutions in this section, we can determine the fluid quantities at the end of the drive beam γ_d , p_d , E_{zd} , and ϕ_d from the drive beam parameters n_d , L_d and the plasma density $n(z)$. We can then calculate φ_{st} and γ_m from the fluid quantities at the end of the drive beam. This map can be used to calculate the wake for any flattop drive beam and any density n . To evaluate this map, we compare the theoretical wake profile with the wake profile computed in 1D PIC simulations for drive beams of various lengths and of densities at, greater than, and less than the plasma densities. As seen in figure 3.4, the theoretical prediction agrees very well with the computed wake.

3.5 Obtaining the phase matched density profile and frequency

We found previously that we could determine a plasma density profile that would guarantee phase matching of a laser pulse and the position of positive density gradient in the

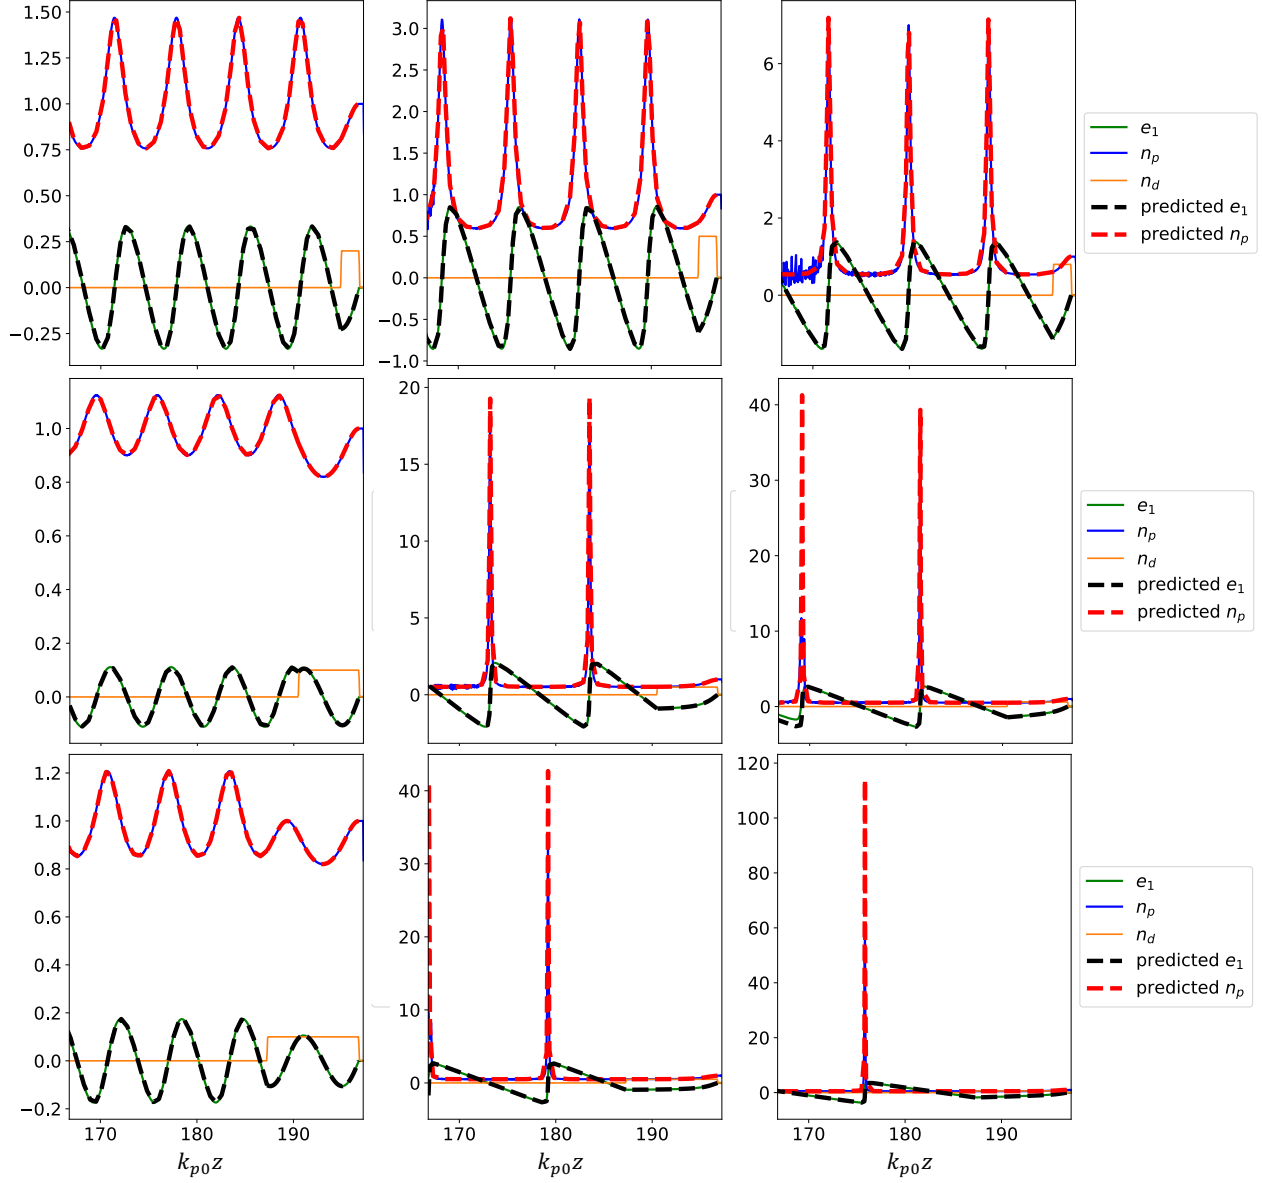


Figure 3.4: Analytic expressions in the wake compared to OSIRIS 1D simulation data. Top row: $k_{p0}L_{d0} = 1.9, n_{d0}/n = 0.2, 0.5, 0.8$. Middle row: $k_{p0}L_{d0} = 6.28, n_{d0}/n = 0.1, 0.5, 0.6$. Bottom row: $k_{p0}L_{d0} = 9.5, n_{d0}/n = 0.1, 0.5, 0.6$. The expressions are exact in a uniform density plasma (here $n = 1$ but any value of n is possible).

wake of an electron beam if we could solve the system of equations (3.4) and (3.7):

$$\frac{d\omega}{dz} = \frac{1}{2k_z c^2} \frac{\partial(\omega_p^2/\gamma)}{\partial\zeta}, \quad (3.75a)$$

$$\frac{dn}{dz} = \frac{1}{2c\omega^2} \frac{\omega_p^2}{\left[\frac{d\zeta_\delta}{dn} \right]^{-1}}. \quad (3.75b)$$

For clarity, we will express the equations in normalized units, with $n \rightarrow n/n_0$, where we use n_0 as a reference plasma density, $\omega_p \rightarrow \omega_p/\omega_{p0}$, and $\omega \rightarrow \omega/\omega_{p0}$ etc., where $\omega_{p0}^2 = e^2 n_0/m_e \epsilon_0$. In the underdense approximation, $k_z c \approx \omega$. Substituting in equation (3.30) for $\frac{\partial(\omega_p^2/\gamma)}{\partial\zeta}$, the frequency shift equation is

$$\frac{d\omega^2}{dz} \simeq n^{3/2} \sqrt{2[\gamma_m - 1]}. \quad (3.76)$$

In dimensionless form, the phase matching condition is

$$\frac{dn}{dz} \simeq \frac{1}{2} \frac{n}{\omega^2} \left[\frac{d\zeta_\delta}{dn} \right]^{-1}. \quad (3.77)$$

Using the work in the previous sections, we have maps from the drive beam parameters n_d , L_d and the unperturbed plasma density n to ζ_δ , the position of positive density gradient where $\delta n = 0$, and $\gamma_m(n; n_{d0}, L_{d0})$, the amplitude of the wake behind the drive beam. From $\zeta_\delta(n; n_{d0}, L_{d0})$, the gradient $\frac{d\zeta_\delta}{dn}(n)$ can be determined numerically. Thus we can close the system of equations (3.76)(3.77) and solve for n , ω . We summarize the procedure for obtaining the density profile and frequency shift in algorithm 1, together with references to the relevant equations in the text. In figure 3.5 we show a sample profile as predicted by algorithm 1. This computation was performed for a drive beam with density $n_d = 0.4$, and length $L_d = 1$. The tapered density profile is shown on the left in (a) and the relative gain in frequency, ω/ω_0 , is show on the right in (b).




Figure 3.5: Numerically determined profiles for unlimited photon acceleration, computed by algorithm 1 with $n_d = 0.4$, $L_d = 1$, and $A = 0.4$. (a) shows the computed plasma density and (b) shows the relative gain in laser frequency.

Algorithm 1 Determining the PA^∞ density profile and expected frequency shift

- 1: Input: Drive parameters n_d, L_d
 - 2: Determine x_d implicitly from n_d, L_d ▷ Equations (3.49), (3.54), (3.67)
 - 3: Evaluate γ_d, p_d, E_{zd} , and ϕ_d as functions of x_d ▷ Equations (3.36)
 - 4: Evaluate φ_{st} and γ_m as functions of γ_d, p_d, E_{zd} , and ϕ_d ▷ Equations (3.22), (3.74)
 - 5: Use $\zeta_d(n)$ to numerically evaluate $\frac{\partial \zeta_d}{\partial n}$ ▷ Equation (3.20)
 - 6: Numerically solve $\frac{d\omega^2}{dz} \simeq n^{3/2} \sqrt{2[\gamma_m - 1]}$, $\frac{dn}{dz} \simeq \frac{1}{2} \frac{n}{\omega^2} \left[\frac{d\zeta_d}{dn} \right]^{-1}$ ▷ Equations (3.76)(3.77)
-

Eqns. (3.76) and (3.77) are a coupled system of differential equations that can be evaluated to determine the density profile and expected frequency shift for phase matched photon acceleration driven by an ultrarelativistic beam driver. However, finding the variation of ζ_δ with density in the nonlinear case is a challenge as it is the solution of an implicit equation. This can be evaluated numerically for specific cases but is not tractable for analytic predictions of long-term behavior or scaling laws. Insight can be gained by using ultrashort and weak driver approximations, for which analytic solutions can be found. The ultrashort and weak driver limits are presented in the following section.

3.6 Ultrashort and weak driver limits and unlimited photon acceleration

In this section, we consider the phase matched photon acceleration in the ultrashort and weak driver limits. Analytic expressions are obtained for the frequency ω and unperturbed plasma density n . Following the asymptotic trajectory of the ω and n as propagation distance z tends to ∞ , it is demonstrated that $\omega \rightarrow \infty$. This mathematically rigorous, arbitrary frequency shift motivates the phrase ‘unlimited photon acceleration.’

3.6.1 Ultrashort limit

The ultrashort approximation means fixing $A = n_d L_d$ and letting $L_d \rightarrow 0$, so that $n_d = A/L_d \rightarrow \infty$. Recall Eqn. (3.74), $\gamma_m = 1 - n_d \phi_d$, and that in the case that the driver density n_d is larger than twice the background density, $n_d > 2n(z)$, the electric potential ϕ_d satisfies the relation $\phi_d = -\tan^2 \theta_d$ and θ_d satisfies Eqn. (3.67),

$$(2n_d/n - 1)k_p L_d = 2\sqrt{2n_d/n} \left(k'^2 F(\theta_d, k) - E(\theta_d, k) + \tan \theta_d \sqrt{1 - k^2 \sin^2 \theta_d} \right), \quad (3.78)$$

where $k^2 = n/2n_d$.

Evaluating in the ultrashort limit, where $\theta_d \ll 1$, we get

$$\begin{aligned} \frac{(n_d/n)k_p L_d - k_p L_d/2}{\sqrt{2n_d/n}} &\simeq (1 - k^2) \left[\left(\theta_d + \frac{k^2 \theta_d^2}{8} \right) - \left(\theta_d - \frac{k^2 \theta_d^2}{8} \right) \right] + \theta_d \left[1 - \frac{k^2}{2} (1 + \theta_d^2) \right] \\ \frac{A}{\sqrt{n}} \frac{\sqrt{n}}{\sqrt{2n_d}} &\simeq \theta_d - k^2 \theta_d + \frac{k^2 \theta_d^2}{8} - \theta_d + \frac{k^2 \theta_d^2}{8} + \theta_d \left[1 - \frac{k^2}{2} + \left(-\frac{k^2}{2} \right) \frac{\theta_d^2}{8} \right] \\ \frac{A}{\sqrt{2n_d}} &\simeq \theta_d. \end{aligned} \quad (3.79)$$

This leads to expansions for the potential ϕ_d at the end of the drive beam,

$$\phi_d \simeq -\frac{A^2}{2n_d}, \quad (3.80)$$

and the amplitude γ_m of the wake behind the driver,

$$\gamma_m \simeq 1 + \left(\frac{n_d}{n}\right) \left(\frac{A^2}{2n_d}\right) = 1 + \frac{A^2}{2n}. \quad (3.81)$$

Hence, the frequency shift relation in the ultrashort limit is

$$\frac{d\omega^2}{dz} \simeq nA. \quad (3.82)$$

Eqns. (3.82) and (3.77) may be combined,

$$\begin{aligned} \frac{d\omega^2}{dz} &= \frac{d\omega^2}{dn} \frac{dn}{dz} = nA \left(\frac{1}{2\omega^2} \left[\frac{d\zeta_\delta}{dn} \right]^{-1} \right)^{-1} \\ &= 2A\omega^2 \frac{d\zeta_\delta}{dn}, \end{aligned} \quad (3.83)$$

and directly integrated,

$$\int_{\omega_0}^{\omega} \frac{1}{\omega'^2} d\omega' = 2A \int_{n_0}^n \frac{d\zeta_\delta}{dn'} dn', \quad (3.84)$$

to obtain

$$\omega(n) = \omega_0 \exp [A (\zeta_\delta(n) - \zeta_{\delta 0})], \quad (3.85)$$

where $n(0) = n_0$. Since $d\omega^2/dz > 0$ and $dn/dz < 0$, there are no fixed points or periodic orbits and so, by the Poincaré-Bendixson theorem [105], there are no limit sets to the orbits.

In fact, it can be shown that as $n \rightarrow 0$, $\zeta_\delta(n) \rightarrow \infty$ and therefore $\omega(n) \rightarrow \infty$.

3.6.2 Unlimited photon acceleration

In this section we show that photon acceleration is a function of propagation distance within the model defined by the phase matching condition, Eqn. (3.77), and the ultrashort frequency shift relation, Eqn. (3.82), or as expressed in the integrated relation for frequency $\omega(n)$ as a function of n given by Eqn. 3.85. That is, arbitrary frequency shifts can be achieved provided the wake can be sustained for a prescribed distance and therefore the photon

acceleration is potentially unlimited. There are several steps involved in the demonstration. First, φ_{st} is shown to be $\pi/2$. Then, a simpler expression for ζ_δ is found. This expression is then shown to diverge as $n \rightarrow 0$.

First we determine $\varphi_{st} = \pi/2$. Using the ultrashort approximation $\gamma_m \simeq 1 + A^2/2n$, we can find expansions for γ_d , E_{zd} , and p_d and combine them to determine $\varphi_{\text{wakestart}} = \tan^{-1} \left(-\frac{E_{zd}}{p_d} \sqrt{\frac{1+\gamma_d}{2}} \right)$. Evaluating E_{zd} :

$$\begin{aligned} E_{zd} &= \sin 2\theta_d \sqrt{2(1 - n_d/n) - 1/x_d + (2n_d/n - 1)x_d} \\ &\simeq +A\sqrt{n_d}/n. \end{aligned} \quad (3.86)$$

Now we evaluate p_d :

$$p_d = \frac{1}{2} \left(\frac{1}{x_d} - x_d \right) \simeq -\frac{A^2}{4n}. \quad (3.87)$$

Finally, we expand γ_d :

$$\gamma_d = \frac{1}{2} \left(\frac{1}{x_d} + x_d \right) \simeq \frac{1}{2} \left(1 - \frac{A^2}{2n} + 1 + \frac{A^2}{2n} \right) = 1. \quad (3.88)$$

Combining the three, we find the expansion for φ_{st} :

$$\tan \varphi_{\text{wakestart}} \simeq -\frac{A\sqrt{n_d}/n}{-A^2/4n} \sqrt{\frac{1+1}{2}}, = \frac{n_d}{A} \rightarrow \infty \quad (3.89)$$

and hence $\varphi_{st} \rightarrow \pi/2$ as $L_d \rightarrow 0$.

Next an ultrashort expression for ζ_δ is derived. Note that in the ultrashort limit, equations (3.19) and (3.20) are

$$\begin{aligned} \zeta_\delta &= \zeta(3\pi/2) - \zeta(\pi/2) \\ &= \frac{1}{\omega_p} \left[\frac{2}{\kappa'} \left\{ E \left(\frac{3\pi}{2}, \kappa^2 \right) - E \left(\frac{\pi}{2}, \kappa^2 \right) \right\} - \kappa' \left\{ F \left(\frac{3\pi}{2}, \kappa^2 \right) - F \left(\frac{\pi}{2}, \kappa^2 \right) \right\} + 4 \frac{\kappa}{\kappa'} \right]. \end{aligned} \quad (3.90)$$

Using the identities $E(3\pi/2, \kappa) = 3E(\pi/2, \kappa)$ and $F(3\pi/2, \kappa) = F(\pi/2, \kappa)$ [104], ζ_δ further

simplifies to

$$\zeta_\delta = \frac{1}{\omega_p} \left[\frac{4}{\kappa'} E(\kappa) - 2\kappa' K(\kappa) + 4\frac{\kappa}{\kappa'} \right]. \quad (3.91)$$

Now the ultrashort expression ζ_δ can be evaluated in the limit $n \rightarrow 0$. Recall that $\kappa^2 = \frac{A^2}{4n+A^2} \rightarrow 1$ as $n \rightarrow 0$ and so $\kappa'^2 \rightarrow 0$. It can be shown [104] that $E(\kappa) \rightarrow 1$ as $\kappa \rightarrow 1$ and

$$K(\kappa) \sim \ln \left(\frac{4}{\kappa'} \right) \quad \text{as } \kappa \rightarrow 1. \quad (3.92)$$

Now ζ_δ simplifies to

$$\zeta_\delta \sim \frac{1}{\sqrt{n}} \left[\frac{4}{\kappa'} (1+1) - 2\kappa' \ln \frac{4}{\kappa'} \right]. \quad (3.93)$$

Because $\lim_{x \rightarrow 0} x \ln x = 0$,

$$\zeta_\delta \sim \frac{1}{\sqrt{n}} \frac{8}{\kappa'} \rightarrow \infty \quad \text{as } n \rightarrow 0. \quad (3.94)$$

Hence, the laser frequency ω is unbounded as plasma density tends to 0 and the model name ‘unlimited photon acceleration’ is justified. This is quite significant in that a physical mechanism for frequency shift is now rigorously proven to admit arbitrary frequency shifts, provided that the drive beam and wake can be sustained.

Larger frequency shifts require lower densities and longer propagation distances. For a given frequency ω to be realized, equation (3.85) indicates the plasma density n that must be attained. Then, the propagation distance z can be determined as a function of n . The differential relation between z and n is

$$\frac{dz}{dn} = \frac{d\omega^2}{dn} \frac{dz}{d\omega^2} = \frac{2\omega^2}{n} \frac{d\zeta_\delta}{dn}. \quad (3.95)$$

This equation can be integrated to give

$$z = 2\omega_0^2 \int_{n_0}^n \frac{1}{n'} \frac{d\zeta_\delta}{dn'} e^{2A(\zeta_\delta(n) - \zeta_{\delta 0})} dn', \quad (3.96)$$

which implicitly defines the phase-matched density profile $n(z)$. Equations (3.85) and (3.96) provide a complementary and more analytic method of determining density profiles and frequency shifts to the numerical integration of equations (3.76) and (3.77).

3.6.3 Weak driver limit

For a moderate strength driver, $A < 1$, we can use small angle expansions to approximate the quantities

$$\kappa^2 = \frac{\gamma_m - 1}{\gamma_m + 1} = \frac{A^2/2n}{2 + A^2/2n} = \frac{A^2}{4n} - \frac{A^4}{16n^2} + \dots, \quad (3.97a)$$

$$\kappa'^2 = \frac{2}{\gamma_m + 1} = \frac{1}{1 + A^2/4n} = 1 - \frac{A^2}{4n} + \frac{A^4}{16n^2} + \dots, \quad (3.97b)$$

$$E(\varphi, \kappa) = \int_0^\varphi \sqrt{1 - \kappa^2 \sin^2 \varphi'} d\varphi' = \varphi(1 - A^2/16n) + \sin(2\varphi)A^2/32n + \mathcal{O}(A^4), \quad (3.97c)$$

$$F(\varphi, \kappa) = \int_0^\varphi 1/\sqrt{1 - \kappa^2 \sin^2 \varphi'} d\varphi' = \varphi(1 + A^2/16n) - \sin(2\varphi)A^2/32n + \mathcal{O}(A^4). \quad (3.97d)$$

The function $\zeta(n)$ is

$$\omega_p \zeta = \varphi - \frac{A}{n^{1/2}} \sin \varphi + \frac{A^2}{16n} (3\varphi + \sin 2\varphi). \quad (3.98)$$

The location of the zero density perturbation, ζ_δ , can be described accurately by the expansion

$$\zeta_\delta = \frac{\pi}{n^{1/2}} + \frac{2A}{n} + \frac{3\pi}{16} \frac{A^2}{n^{3/2}} + \dots. \quad (3.99)$$

Retaining up to the third term in this expansion, the phase-matching relation, Eqn. 3.77, is

$$\frac{dn}{dz} = -\frac{n^{5/2}}{\pi\omega^2} \left(1 + \frac{4A}{\pi n^{1/2}} + \frac{9A^2}{16n} \right)^{-1}. \quad (3.100)$$

The expansions developed in this section are seen to be accurate over a larger parameter range than may be expected. In figure 3.6 we present γ_m , ζ_d , n , and k_L for $A = 0.01, 0.1, 1$, and 10 and for $k_{p0}L_{d0} = 0.01$. This short driver length guarantees that the ultrashort limit holds very well. This is why the dashed and solid lines in the top row, which shows the

wake amplitude γ_m , agree so well for all values of A calculated. The second row shows the position of zero density perturbation ζ_δ and its approximations, which agree with the exact function to $A = 1$. The next rows show the predicted density profile and frequency gain. The density profile is accurately predicted with the expansion for A as large as 1. The slow variation in n in all cases means that the predicted frequency gain of the analytic expansion matches the numerical prediction for all A . We see that generally the expansion is good for $A = 1$ and surprisingly accurate even for $A = 10$.

3.7 3D considerations

There are several multi-dimensional effects to consider in using the 1D density profile for frequency upshift. These include laser pulse diffraction, the transverse wake profile, drive beam self-focusing and the qualitatively different axial density profile in the blowout regime. These affect both the laser pulse dynamics, the electron beam dynamics, and the wake profile.

In one dimension, the region of frequency upshift has length $\leq \lambda_p/2$ since there is no transverse evolution and so for the entire laser pulse to see frequency upshift, the pulse width must be shorter than $\lambda_p/2$. We note, as did Esarey et al. [17] that the wake itself has the shape of a plasma channel and can focus a laser pulse. The region of frequency upshift and laser focusing has width $\lambda_p/4$; for frequency upshift and guiding, the pulse width must be less than $\lambda_p/4$. Additionally, for higher intensity pulses above the critical power, relativistic self-focusing can maintain the pulse profile.

In linear theory, the wake profile is the product of the longitudinal and transverse profiles, where the transverse profile of the wake is the transverse profile of the drive beam. Given a drive beam creating a quasilinear wake with a transverse gaussian profile, the wake transverse profile will be gaussian. On axis, the laser pulse experiences a parabolic transverse density profile, which if deep and wide enough will guide the laser pulse. We use a gaussian drive

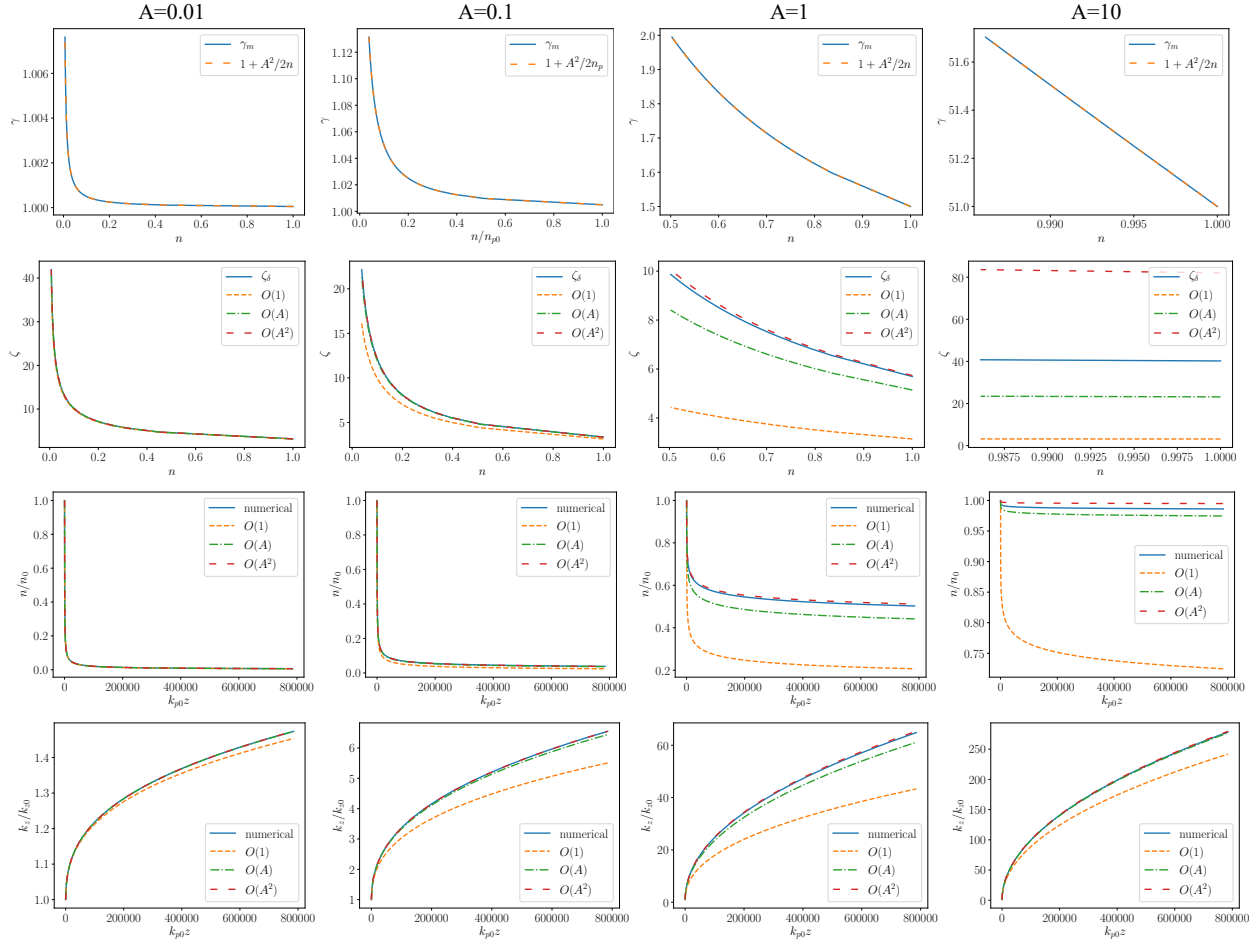


Figure 3.6: Comparing numerical and ultrashort small A expansions. We show that the wake amplitude γ_m , the position of 0 density perturbation ζ_δ , and the predicted density n and frequency gain k_L/k_{L0} profiles for $A = 0.01, 0.1, 1$, and 10 when $k_{p0}L_{d0} = 0.01$. The length satisfies $k_{p0}L_d \ll 1$ so $\gamma_m = 1 + A^2/2n$ is always accurate. The expansion to $\mathcal{O}(A^2)$ is quite accurate for ζ_δ, n , and k_L for A as large as 1 . The expansion when $A = 10$ is not good for ζ_δ but is surprisingly accurate in predicting n and k_L even for $A = 10$.

electron beam and find that the laser pulse stays focused over more than 100 Rayleigh ranges.

The tail of the electron drive beam experiences radial focusing forces in the wake created by the drive beam and is pinched. Once the beam radius decreases below a threshold value, the wake develops the quasi-spherical shape characteristic of the blowout regime [5, 106], a regime unique to 2- and 3-dimensions. In the blowout regime, the longitudinal wake profile becomes singular and the region with a density gradient becomes vanishingly small with slope tending to infinity. The laser pulse doesn't fit in the region of frequency upshifting and so the blowout wake regime is to be avoided. Even if the drive beam is initially broad and creates an almost 1D wake, in simulations the beam can be pinched until it creates a wake in the blowout regime. This is a significant effect that is detrimental to the current frequency shifting scheme and we are actively investigating how to mitigate this.

3.8 Results

We use the profile determined by equations (3.76) and (3.77) in a series of simulations to demonstrate unlimited photon acceleration.

- First, we demonstrate the utility of the phase matched profile by comparing a simulation of a uniform plasma profile with the phase matched PA^∞ profile. The laser pulse quickly dephases from the upshifting region of the wake.
- Next we use realistic beam parameters to demonstrate $10\times$ frequency shift. This simulation will then be repeated in 3D; comparison indicates the effective 1D nature of the scheme even in 3D geometry.
- We perform a 1D simulation with a long drive beam to demonstrate a profile not obtainable in the ultrashort limit. This simulation demonstrates $30\times$ frequency shift and $6\times$ energy gain

- We repeat the previous simulation with a piecewise approximation to the density profile and still realize $30\times$ frequency shift.
- We perform a quasi-3D simulation demonstrating $10\times$ frequency shift and $5\times$ energy gain
- Finally, we demonstrate $10\times$ frequency shift and $5\times$ energy gain in quasi-3D using a piecewise approximation to the phase matched profile. This demonstrates the robustness of the phase matched profile to perturbations.

The simulations in this chapter are performed using particle-in-cell (PIC) codes [27, 28]. The particle-in-cell method represents the plasma phase-space with finite particles. To calculate forces on the particles, charge and current densities are evaluated on a grid by interpolating charge and current from the particle positions to the grid. Then Maxwell's equations are solved on the grid to obtain the electric and magnetic fields. Then the fields are interpolated back onto the particle positions and forces on the particles can be determined. The particle equations of motion are evaluated and the particle positions and momenta are updated. This completes a cycle called the PIC cycle for updating the particle positions and electromagnetic fields, as shown in figure 3.7. The simulations in this dissertation use the 1D version of the PIC code OSIRIS [107, 108] and the quasi-3D code FBPIC [109].

To measure the efficiency of the PA^∞ scheme, we define the average energy transfer efficiency as

$$\eta = \frac{\Delta U_{laser}}{-\Delta U_{beam}} = \frac{U_{laser} - U_{laser,initial}}{U_{beam,initial} - U_{beam}}, \quad (3.101)$$

with energy changes measured between the initial and final times.

3.8.1 1D simulations with and without the PA^∞ density profile

In this section we use 1D simulations to compare position of the laser in the case of the tapered plasma density profile with a uniform plasma density case. The drive beam has




Figure 3.7: Schematic of the PIC cycle.

density $n_d = 0.4n_{p0}$ and length $k_{p0}L_d = 1.0$ for $A = 0.4$. Initially the laser has $\omega_L = 10\omega_{p0}$. The laser has a normalized vector potential a_0 of 0.5, with an intensity full-width at half-max (FWHM) of 3 fs. We start with 618 cells per λ and in the tailored density case end with 61.8 cells per λ . We use a CFL condition of $\Delta t = 0.9995\Delta x$ and 4 particles per cell.

We see in figure 3.8 the position of the laser centroid in the wake, as a function of propagation distance. Three panels are shown, comparing the effects of a uniform plasma density with the PA^∞ tapered plasma density profile and a piecewise-constant approximation to the PA^∞ density profile. The top panel shows the laser centroid and wake in the case of uniform plasma density. In the uniform density case, the laser centroid clearly slips past the density peak and out of the frequency-shifting region, back into the density gradient of the next density peak, past that and on through the wake. No sustained frequency shift is attained. On the other hand, the middle panel shows the benefit of the PA^∞ tapered plasma density profile. The laser centroid maintains its position in front of the first density peak in the wake, at the phase where the density gradient is increasing with $\tau = -c\xi$ and the frequency shifts up. The bottom panel shows the effect of a piecewise-constant approximation to the PA^∞ plasma density profile. The wake shows discontinuities at the plasma density

jumps, but the laser centroid maintains a smooth phase in the region of increasing density gradient, near where $\delta n = 0$. Even an approximation to the PA^∞ density profile can be clearly seen to eliminate dephasing of the laser in the wake.

3.8.2 1D PA^∞ simulation for comparison with quasi-3D simulation

In this section we perform a 1D simulation using the PA^∞ density profile and drive beam parameters also used in a later quasi-3D simulation. Initially the laser has $\omega_L = 10 \omega_{p0}$. The drive beam has density $n_d = 0.2 n_{p0}$ and length $k_{p0}L_d = 1.9$. These parameters correspond to $A = 0.38$, which is the value of A used in the quasi-3D simulation. The laser has a normalized strength parameter $a_0 = \frac{eE}{m\omega_{p0}}$ of 2.4, with an intensity full-width at half-max (FWHM) of 3 fs. We start with 618 cells per λ and in the tailored density case end with 61.8 cells per λ . We use a CFL condition of $\Delta t = 0.9995\Delta x$ and 4 particles per cell.

We can see in figure 3.9 that the mean laser wavenumber matches the predicted wavenumber and increases by a factor of more than 10. The drive beam is linearly chirped in energy so that each section of the beam is depleted at roughly the same time. The chirp has the trapezoidal form $\gamma(\zeta) = 4000 + 32000(\omega_{p0}\zeta/1.9)$. The form of the chirp is chosen so that most of the drive beam energy is extracted, and 7.2% of that extracted energy is transferred to the witness laser pulse. The frequency shift and energy change stop around $120,000k_{p0}z$, at which point the drive beam is depleted of energy.

3.8.3 1D PA^∞ simulation with a long drive beam

In this section we demonstrate the ability of the model to generate frequency shift for finite length drive beams. We can see shifts of $30\times$ in frequency, $6\times$ in energy, $25\times$ in intensity, and $13\times$ in compression. The frequency gain stops once the driver is depleted of energy, indicating that the gain is limited only by the ability of the driver to maintain the wake.




Figure 3.8: Laser centroid in wake plotted over propagation distance, calculated with and without the PA^∞ tapered plasma density profile. Top panel: uniform density; the laser centroid dephases rapidly from the wake when the plasma density is uniform. Middle panel: PA^∞ tapered density profile; the laser centroid maintains its position relative to the wake. Bottom panel: piecewise-constant approximation to the tapered density profile; the laser maintains phase in the wake even in the case of only a piecewise approximation to the PA^∞ tapered density profile. To illustrate the wake we plot $\delta n = n_w(\xi; z) - n(z)$. Note that in these images the wake coordinate is $\xi = z - ct$ not ζ , but for $v_p \rightarrow c$, $\xi \simeq -c\zeta$.




Figure 3.9: 1D Results for a drive beam with density $n_d = 0.2$ and length $L_d = 1.9$ and a laser pulse with normalized strength parameter $a_0 = 2.4$, (a) spectrum of laser pulse versus propagation distance, the red curve shows the predicted frequency as determined by equations (3.76) and (3.77) and the black line shows the mean frequency, note that the mean frequency increases by a factor of 10, (b) Tapered density profile for dephasingless photon acceleration as determined by equations (3.76) and (3.77), (c) drive beam energy in blue and laser pulse energy in red, shown to different scales, the electron drive beam loses energy linearly to the wake and is almost completely depleted of energy, the laser gains about 7% of the initial energy in the drive beam

The drive beam has density $n_d = 0.3n_{p0}$ and length $k_{p0}L_d = 12.56$. The beam has an up-chirp in energy similar to the previous section of the form $\gamma(\zeta) = 4000 + 32000(\omega_{p0}\zeta/12.56)$. The chirp helps prevent partial beam depletion but is not optimal; an item of future work is better tuning of the chirp. Initially the laser has $\omega_L = 10\omega_{p0}$. The laser has an a_0 of 4.0, with an intensity FWHM of $\omega_{p0}\tau = 0.85$. We use 30000 cells for 618 cells per λ_{L0} and about 21 cells per λ_{Lf} , a CFL condition of $\Delta t = 0.9995\Delta x$, and 4 particles per cell.

Figure 3.10 shows that the mean laser wavenumber matches or exceeds the predicted wavenumber and increases by a factor of more than 30. About 5% of the energy lost by the drive beam is transferred to the witness laser pulse. The frequency energy transfer sees a discontinuity around $60,000k_{p0}z$. This is the drive beam depletion length; at this point a portion of the beam has lost enough energy to move at sub-relativistic velocities and gets caught by the electron-accelerating fields. This loads the wake, so the laser pulse behind no longer sees an increasing density gradient and doesn't experience any more frequency shift.

The reported values of this 1D dimensionless simulation can be understood in physical units if we impose some initial scales. Assuming the initial laser wavelength is 800 nm and




Figure 3.10: 1D results for a drive beam having density $n_d = 0.3$, length $L_d = 12.56$ and a laser pulse having normalized strength parameter $a_0 = 4.5$, (a) spectrum of laser pulse versus propagation distance, the red curve shows the predicted frequency as determined by equations (3.76) and (3.77) and the black line shows the mean frequency, note that the mean frequency increases by a factor of 30, (b) Tapered density profile for dephasingless photon acceleration as determined by equations (3.76) and (3.77), (c) drive beam energy in blue and laser pulse energy in red, shown to different scales, the electron drive beam loses energy linearly to the wake and loses about half its energy before parts of the beam are depleted and the wake is disrupted. Laser energy increases monotonically, with most of the energy gain occurring early in the simulation.

the drive beam has an initial root-mean-square radial extent $\sigma_r = 16 \mu\text{m}$, then the system propagates about 10 cm and the shifts are $30\times$ in frequency, from an initial wavelength $\lambda = 800 \text{ nm}$ to 27 nm , $6\times$ in energy, from $U = 50 \text{ mJ}$ to 900 mJ , $25\times$ in intensity, and $13\times$ in compression, from an initial full-width half-maximum in intensity pulse duration of $c\tau = 900 \text{ nm}$ to 69 nm .

3.8.4 1D PA^∞ with a piecewise approximation to the exact profile

The PA^∞ plasma density profile is integral to sustained frequency shift, so an important consideration is how sensitive the shifts are to the profile. To investigate the robustness of PA^∞ , we repeat the previous simulation using a piecewise constant approximation to the tapered plasma density profile. We still measure $30\times$ frequency shift and $6\times$ energy gain, demonstrating that the density profile is robust to perturbations.

The drive beam has density $n_d = 0.3n_{p0}$ and length $k_{p0}L_d = 12.56$. Initially the laser has $\omega_L = 10\omega_{p0}$. The laser has an a_0 of 4.5, with an intensity FWHM of $3fs$. We use 30000 cells for 618 cells per λ_{L0} and about 21 cells per λ_{Lf} , a CFL condition of $\Delta t = 0.9995\Delta x$, 4

particles per cell.

Figure 3.11 shows that the mean laser wavenumber matches or exceeds the predicted wavenumber and increases by a factor of more than 30. About 5% of the energy lost by the drive beam is transferred to the witness laser pulse. The frequency energy transfer sees a discontinuity around $60,000k_{p0}z$. As in the previous section, a portion of the beam is depleted of energy, dephases, loads the wake, and stops the frequency shift of the trailing laser pulse.




Figure 3.11: Results in 1D for drive beam $n_d = 0.3$, $L_d = 12.56$ and laser $a_0 = 4.5$, (a) spectrum of laser pulse versus propagation distance, measured in k_L/k_{L0} where $k_L = k_z$ is the laser wavenumber, the red curve shows the predicted frequency as determined by equations (3.76) and (3.77) and the black line shows the mean frequency, note that mean frequency increases by a factor of 30, (b) Tapered density profile (dashed red) for dephasingless photon acceleration as determined by equations (3.76) and (3.77) and piecewise approximation (solid blue), (c) drive beam energy in blue and laser pulse energy in red, shown to different scales, the electron drive beam loses energy linearly to the wake and loses about half its energy before parts of the beam are depleted and the wake is disrupted.

This simulation uses a piecewise approximation to the density profile obtained from solving equations (3.76) and (3.77). Using a scaling for $100\times$ underdense initial plasma density n_0 relative to an initial laser wavelength of $\lambda_{L0} = 800$ nm, the step lengths are 1, 1, 4, 6, 8, 10, and 70 mm. The density over each interval is the density of the numerically evaluated tapered profile at the midpoint of the interval. Significant frequency shift is still obtained, demonstrating that the phase matched photon acceleration scheme is robust to perturbations in the density profile, and could be achieved with a sequence of gas cell targets of increasing length and decreasing plasma density.

3.8.5 Quasi-3D PA^∞ simulation with a short, broad drive beam

We use the PIC code FBPIC [109] to study PA^∞ in a quasi-3D geometry. FBPIC runs on GPU. The advantage to FBPIC is its spectral cylindrical representation. The spectral field solve reduces numerical dispersion with fewer grid points than the standard Yee finite difference scheme. The spectral cylindrical representation decomposes the azimuthal direction into a Fourier representation. Our problem is nearly cylindrically symmetric, with the most significant asymmetry in the linear polarization of the laser. We use two angular modes, which is enough to capture the linearly polarized laser as well as the cylindrically symmetric wake from the drive beam.




Figure 3.12: Isocontours from quasi-3d simulation data demonstrating phase matched photon acceleration; the purple disk at the right is the drive electron beam, yellow and green are the positive and negative density perturbations of the wake, and the blue and red spheroid is a contour of the envelope of the transverse electric field, colored to show the laser pulse field phase.

We use the quasi-3d particle-in-cell code FBPIC [109] to demonstrate unlimited photon acceleration. In the simulation, the drive electron beam has a charge $Q = 8.6$ pC, length $c\tau = 0.4 \mu m$, focused radial extent $\sigma_r = 32 \mu m$, and a divergence $\sigma_{r'} = 0.4$ mrad. The beam has a linear up-chirp of the form $\gamma(\zeta) = 5 \times 10^4 + 5 \times 10^4(c\zeta/L_d)$ for an average

beam particle energy of 50 GeV. The chirp mitigates beam erosion somewhat but there is room for tuning. The beam parameters are similar to those accessible at the SLAC FACET-II facility [110]. The laser has a normalized field strength parameter $a_0 = 1$ and initial wavelength $\lambda_{L0} = 800$ nm, corresponding to a peak intensity $I = 2.1 \times 10^{18}$ Wcm⁻², a pulse width $c\tau = 1.25$ μ m, and a spot size of $w_0 = 2\lambda_{p0} = 20\lambda_{L0} = 16$ μ m, starting with its centroid $\zeta_{\text{cent}} = 6.6$ μ m/c behind the front of the electron drive beam. For these parameters, $\zeta_\delta = 5.4$ μ m/c, so we see some robustness to variation in initial position of the laser pulse. Single cycle laser pulses at this intensity have been demonstrated experimentally [97]. Visualization of the drive beam, its wake and the witness laser pulse in the quasi-3d simulation is shown in Fig. 3.12.

We use the plasma density profile shown in Fig. 3.13 and determined through algorithm 1 from equations (3.76- 3.77), scaled by $n_0 = 1.74 \times 10^{19}$ cm⁻³, which is 100 \times less than the initial critical density of the laser. For numerical parameters we use 3200×200 grid points and 2 angular modes. This gives 85 cells per λ_{L0} and about 8 cells per λ_{Lf} , which therefore requires a dispersion-free solver or numerical dispersion would be an impediment. For the plasma we use 2×2 particles per cell (ppc) in the z and r directions times 8 ppc in the azimuthal direction.

The results of the quasi-3d simulation of PA ^{∞} can be seen in Fig. 3.14. Figure (a) shows the spectrum as a function of length propagated in the simulation. Overlaid is the theoretical model of Eqns. (3.76) and (3.77) (red line, identical to the analytic predictions of equations (3.85 and 3.96)), which predicts the shift of the photons well. Figure (b) shows the initial spectrum and the spectrum at the end of the simulation. As well as showing that the pulse maintains a narrow bandwidth, the amplitude of the spectrum has increased, indicating the energy of the laser pulse increased. Figure (c) shows the energy evolution of the drive beam and the laser pulse. The drive beam loses 25 J of energy over the course of the simulation




Figure 3.13: Density profile for unlimited photon acceleration driven by an ultrashort driver with $A = (k_{p0}L_{d0})(n_{d0}/n_0) = 0.38$.

while the laser pulse, initially having 50 mJ, gains 200 mJ for about 1% energy transfer efficiency, which could be increased by optimization of the driver. Note that overall pulse energy gain is not unexpected, as to within the quasistatic approximation, local field action is conserved [15], and so the energy gain by the pulse would be expected to scale with the frequency increase. In the simulations, the nearly $10\times$ frequency shift results in a lower $5\times$ pulse energy gain, because of losses due to diffraction/dispersion.

Due to the nonlinear wake, the back of the pulse sees a steeper plasma gradient than the front of the pulse. The pulse develops significant up-chirp, which, together with dispersion, leads to a $2.7\times$ compression of the pulse from an initial duration of $c\tau = 1.25 \mu\text{m}$ to $0.51 \mu\text{m}$. The intensity sees a significant increase of $20\times$ from an initial intensity $I = 2 \times 10^{18} \text{ W/cm}^2$ to $4 \times 10^{19} \text{ W/cm}^2$, a relativistic intensity corresponding to $a_0 = 0.4$ at 80 nm.

We note that the laser pulse stays focused throughout the course of the simulation. The drive beam has a gaussian radial profile, which is therefore parabolic close to axis. Thus the wake has an approximately parabolic transverse profile and acts as a guiding channel for the laser pulse. The pulse stays focused over more than 10 centimeters when the Rayleigh range corresponding to the final wavelength is $z_R = \pi w_0^2/\lambda_{Lf} \approx 10 \text{ mm}$, meaning that the pulse




Figure 3.14: Quasi-3D simulation results from FBPIC for a drive beam with $k_{p0}L_{d0} = 0.5$. (a) Comparing the initial and final spectra. The spectrum has shifted to about $10\times$ the initial wavenumber. (b) Plot of the spectrum over time. The red line is the theoretical 1D frequency shift. (c) Drive beam energy (blue) and witness laser energy (red)

was effectively guided for more than 10 Rayleigh lengths.

The simulation is shown ending at 72 mm. Beyond this length, the beam driver has lost sufficient energy to break up. Going beyond this length would also start to reach the resolution limit of the simulation for the maximum frequency. In 1d simulations at higher resolution with a longer drive pulse, where tuning of the drive beam propagation was easier through parametric studies of many simulations, we have demonstrated even $30\times$ frequency shifts. Beginning with the 2nd harmonic of a Ti:Sapph laser at 400 nm, a 30 times shift (scaling the plasma / initial conditions to maintain similarity and fixing a_0) would result in 125 mJ pulses of coherent 13.3 nm light, which may be useful for photolithography.

3.8.6 Quasi-3D PA^∞ simulation with piecewise-density profile

In this section we repeat the simulation of the previous section but use a piecewise-constant approximation to the density profile, as shown in figure 3.15. We still find $10\times$ frequency shift and $5\times$ energy gain, demonstrating that the tapered density profile is robust to variations in plasma density even in 3D.

The approximation is made on progressively longer intervals, as in the previous 1D piecewise simulation. The first two steps are each 1 mm long, then 4 mm, 6 mm, 8 mm, 10 mm, and the last step here is 70 mm long. The density over each interval is the density of the

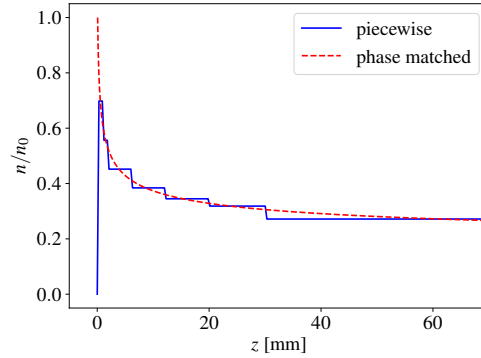


Figure 3.15: Piecewise approximation to the profile used for quasi-3D simulation

numerically evaluated tapered profile at the midpoint of the interval.

We can see the results of propagating through the piecewise constant approximation to the phase matched density profile in figure 3.16. Similar to the simulation with the phase matched density profile, in part (a) we see that the final frequency is almost 10 times the initial frequency, and the amplitude of the spectrum has increased, indicating the energy of the laser pulse increased. The frequency shift is shown over the course of the simulation in (b), following the theoretical prediction indicated by the red line. The center of the distribution follows the theoretical prediction.

In figure 3.14(c) we see the loss of drive beam energy is linear with propagation distance, as the beam energy is continuously transferred to the wake. On the other hand, the laser gains about 200 mJ for a $5\times$ shift in energy. For this simulation about 1% of the energy lost by the drive beam to the wake was transferred to the laser pulse.

Comparing the piecewise results in figure 3.16 with the results using the exact tapered profile 3.14, we see similar frequency shifts and energy gains. We see slightly more shift in the exact case, where the center of the spectrum is above the predicted frequency, than in the piecewise case, where the center of the spectrum is close to the predicted frequency. We also see in the late-time spectra plots in figures 3.16(a) and 3.16(b) that the amplitude of the the spectrum is slightly lower in the piecewise case, and in figures 3.16(c) and 3.16(c)

that the energy gain of the laser pulse is slightly less. These discrepancies are slight and speak strongly to the reliability of the tapered density profile for attaining frequency shift.




Figure 3.16: Quasi-3D simulation results from FBPIC for a drive beam with $k_{p0}L_{d0} = 0.5$ using a piecewise approximation to the phase matched density profile. (a) Comparing the initial and late-time ($z = 72$ mm) spectra. The spectrum has shifted to almost 10x the initial wavenumber. (b) Plot of the spectrum over time. The red line is the theoretical 1D frequency shift. (c) Drive beam energy (blue) and witness laser energy (red)

3.9 Summary

In summary, we have demonstrated a scheme for large frequency upshift of a laser pulse using the wake generated by a relativistic particle beam propagating through a tapered plasma density profile. Our analytic model predicts arbitrary frequency shift limited only by the drive beam, that is, unlimited photon acceleration. We demonstrate the model in several 1- and quasi-3D simulations. These demonstrate how quickly a laser dephases in the wake if the phase-matched density profile is not used. 1D simulations show frequency shifts of $30\times$ or more if a sufficient drive beam is used, as well as demonstrating that drive beams of arbitrary length and density can be used and that non-ideal, physically realizable density profiles still lead to $30\times$ frequency shifts.

In the quasi-3D simulations shown here, a short, unoptimized driver was used, but in principle there is much scope for increasing coupling of driver energy to the wake/witness laser pulse, including advancements in spatiotemporally evolving electron beam drivers [111].

Quasi-3D simulation results directly demonstrated a $10\times$ frequency shift with 1% energy

transfer efficiency. Guiding of the laser pulse is achieved by the wake itself, in addition to the lengthening of the Rayleigh range as the frequency increases and possibly relativistic self-focusing. In quasi-3D simulations we have also demonstrated the robustness of the plasma density profile to perturbations, as well as a route to practical realization of the tapered density profile, using a piecewise constant approximation to the density profile. Frequency shifts of $10\times$ are still observed in quasi-3D simulations with a non-ideal profile.

Our results indicate that using an electron beam driver is a viable route to generate high-power, high-frequency coherent radiation. It may also be possible to demonstrate these shifts with a laser beam driver, although dispersion and frequency shifting of the lead pulse require compensation.

3.9.1 Future work

Several broad directions for future research in unlimited photon acceleration are discussed here. First, understanding 3D transverse effects and feasible acceleration with finite-duration electron beams. Second, optimization of electron beam, laser beam, and plasma density profile. Third, developing the phase matched system of equations for laser-driven photon acceleration. Finally, investigating other methods of maintaining the driver and wake for extended photon acceleration.

The evolution of the drive beam has a significant effect on the wake profile. In particular, the radial focusing forces of the wake cause the drive beam to pinch. A drive beam that is initially broad enough for the 1D wake expressions to be valid will shrink until radial effects change the plasma wavelength and ultimately the highly nonlinear blowout regime is reached, which is undesirable for photon acceleration.

Several open questions remain with regards to understanding and controlling the transverse drive beam dynamics and their effects on photon acceleration. First, while effort was made to avoid the blowout regime in pursuit of finite duration plasma gradients, it may be

possible to realize frequency shifts within the blowout regime. If not, then it is necessary to avoid the blowout regime somehow. Perhaps there are matching conditions so that the electron beam propagates without self-focusing into a blowout-producing profile. It may be that a gaussian beam will self-focus too strongly, but perhaps beams can be constructed having a higher divergence at the back of the beam and so not self-focusing before the depletion length is reached.

Also, simple expressions for the index of refraction and laser group velocity are used. Several nonlinear effects were neglected in the analysis, including relativistic, intensity and channel effects on the wake profile and laser evolution. The effects of these corrections are unknown but could be large enough to affect the wake or laser pulse.

The next broad category of future directions is in optimizing the frequency shift and efficiency of the PA^∞ scheme. As in electron acceleration [112, 113], a properly shaped laser pulse should admit laser beam loading of the wake. This could enhance energy gain, frequency shift, or the amount of light that can be shifted. Also, in this work very short, single cycle laser pulses are used. While such pulses have been experimentally demonstrated [97], the model isn't restricted to single cycle pulses. It is still unknown how well longer pulses can be shifted in this scheme.

The third category of future directions is laser-driven photon acceleration. Laser-driven photon acceleration has been studied more than electron-beam-driven photon acceleration but not with a profile such as was developed in this chapter. In the case of laser-driven photon acceleration a density upramp is used since the frequency-shifting pulse would have a higher group velocity than the drive pulse and so would eventually catch up to the drive pulse. This limits how long and how much the frequency can be increased. There have been theoretical investigations [18] into laser-driven frequency upshift but little successful simulation and even less on dephasingless frequency upshift such as demonstrated in this

chapter.

The last direction of photon acceleration research considered here is concerning the ability of spatiotemporally shaped light and matter [90, 91, 111] to enhance the frequency shifts observed in simulations in this dissertation. We stated previously that ‘unlimited photon acceleration’ is limited by the ability to maintain the wake. Recent research in ‘flying focuses’ and other spatiotemporally shaped light and matter has been applied to electron [89] and photon [21] acceleration. In the previous study of photon acceleration with a ‘flying focus’ [21], a photoionization front was used rather than the finite-extent density gradient and photon acceleration used in this dissertation. Combining a spatiotemporally shaped light or electron beam with the PA^∞ model and density profile introduced in this dissertation could be useful in maintaining the PA^∞ driver over the distances needed for significant frequency shift.

CHAPTER IV

Summary

This dissertation presented two works in the numerical study of plasma waves via kinetic models, resolving the fine-scale structures that arise in phase-space, and the implications for plasma-based acceleration. Chapter I outlined and motivated the dissertation. Chapter II presented the development of an adaptive semi-Lagrangian method called FARSIGHT for the Vlasov-Poisson system inspired by vortex methods in incompressible fluid dynamics. In chapter III the second work was presented, exploring frequency shifting of an electromagnetic wave in a tailored plasma. In the rest of this chapter, we summarize the main results, conclusions, and future directions for each work.

4.1 FARSIGHT

Chapter II presented FARSIGHT, a forward semi-Lagrangian scheme for collisionless electrostatic plasmas described by the 1D1V Vlasov-Poisson equations. The distribution function is represented by quadrilateral panels having a hierarchical tree structure, and each panel is a 3×3 grid of particles which are tracked by 4th order Runge-Kutta timestepping. The electric field is expressed as a convolution integral of the charge density with a regularized electric field kernel. The particles are remeshed at every time step using biquadratic interpolation on the panels, and the panels are adaptively refined to resolve the phase-space

distribution function. The electric field integral is discretized by the trapezoid rule, and the discrete sums are computed by a GPU-accelerated barycentric Lagrange treecode [58].

FARSIGHT performs as well as the state-of-the-art codes in the standard test cases. In calculations of Landau damping, FARSIGHT reproduced the well-known damping and growth rates, while in strong Landau damping, two-stream instability, and beam halo formation, the code was able to resolve small-scale structures in phase-space.

The adaptive mesh refinement was shown to significantly decrease the run time in all cases, especially in the cold two-stream instability and beam halo formation problems. It was shown that conservation errors decrease and fewer negative f -values occur when the numerical parameters are refined.

4.1.1 Future directions

This dissertation has demonstrated the utility of the FARSIGHT method for modeling the Vlasov-Poisson system. However, some of the design choices were made for convenience or simplicity in developing the initial idea and are not optimal. There are several possible directions for future optimization and development of FARSIGHT, falling in the broad categories of algorithmic improvements, additional features for the modeling of additional physics, and extensions of the code to higher-dimensions and different physics regimes.

Several aspects of the FARSIGHT method admit algorithmic and implementation improvements. These possibilities include:

- **Conservative remeshing** Long-time simulations can see significant degradation in physical properties such as positivity, mass, momentum, and energy that should be conserved. The current implementation of FARSIGHT is not exactly conservative but there is work in the literature in limiters for positivity preservation [1] and interpolation schemes that conserve moments [38, 73, 75]. Conservation in the context of AMR is

more involved than these uniform grid methods but some of these ideas could be applied to our adaptive remeshing.

- **Higher order interpolation and quadrature** Finer particle grids in each panel and higher order interpolation and quadrature could admit higher accuracy with fewer panels and reduced computational cost. For example, the current implementation of FARSIGHT could incorporate a biquartic interpolation scheme where each panel consists of 25 points in each direction.
- **More sophisticated refinement scheme** Panel refinement independently in the x and v directions, along with interpolation schemes that are accurate on panels with extreme aspect ratios. This could allow more efficient simulation of problems such as Landau damping, where there is filamentation in the v direction suggesting a much higher need of resolution in v than in x .
- **Improved treecode** The current treecode incorporated in FARSIGHT, the Barycentric Lagrange Treecode (BLTC), is $\mathcal{O}(N \log N)$ and is designed for 3D geometry. There is an $\mathcal{O}(N)$ improved version of the BLTC called the Barycentric Lagrange Dual Tree Traversal (BLD TT) fast multipole method. Speedups are anticipated in using a 1D version of the BLD TT.

Some additional physics are accessible through features that can be added without significantly altering the current FARSIGHT code. Examples include:

- **An external driver** This would allow for the study of driven waves and dynamically accessible phase-space structures [24, 114, 115, 116].
- **Relativistic time stepping** This would allow for the basic test problems of Landau damping and the two-stream instability in the relativistic regime [117]. Relativistic

time integration is also prerequisite to many applications of interest that require other simple features, as will be explained subsequently.

- **Tracking of multiple plasma species** This allows for ion-electron or positron-electron problems in basic physics [118] and more applied problems such as ion shock acceleration [12, 13].

We believe the approach developed here can be extended to higher-dimensional electromagnetic plasmas in confined geometry, using an integral formulation of Maxwell’s equations. The adaptive refinement of phase-space should be even more beneficial for higher dimensional implementations of FARSIGHT and the use of the treecodes BLTC or BLDTT means the method already has access to a fast summation method that scales well on multiple GPUs [77]. These extensions open up challenging problems such as the study of energy transport in hohlraums for inertial confinement fusion or in tokamaks for magnetic confinement fusions, or potentially 3D simulations of plasma-based acceleration.

4.2 Unlimited photon acceleration

Chapter III presented unlimited photon acceleration (PA^∞), a scheme for dephasingless photon acceleration in a particle-beam-driven wake. This scheme is based on the phenomenon that electromagnetic radiation seeing a decreasing plasma gradient shifts up in frequency. In PA^∞ , a laser pulse is situated in the wake behind a relativistic electron bunch so that, in moving window coordinates, it sees an increasing plasma density gradient. An appropriately tapered plasma density (meaning plasma density that decreases in propagation distance) allows the witness laser pulse to stay where the plasma density is increasing in the moving window coordinate. Shifts of $30\times$ in frequency ($\lambda : 800 \text{ nm} \rightarrow 25 \text{ nm}$), $6\times$ in energy ($50 \text{ mJ} \rightarrow 900 \text{ mJ}$), $25\times$ in intensity, and $13\times$ in compression (pulse duration: $3 \text{ fs} \rightarrow 230$ as were observed in 1D simulations. Shifts of $10\times$ in frequency ($800 \rightarrow 80 \text{ nm}$), $5\times$ in energy

(50 \rightarrow 250 mJ), 20 \times in intensity ($2 \times 10^{18} \rightarrow 4 \times 10^{19}$ W/cm²), and 2.7 \times in compression (pulse duration 4.9 fs \rightarrow 1.7 attoseconds(as)) were observed in quasi-3D simulations.

4.2.1 Future directions

As discussed in the summary of chapter III, there are several promising directions for future research in beam-driven photon acceleration.

- **understanding 3D transverse effects in finite-duration electron beams**

The studies in this dissertation used a (a) longitudinally short, (b) radially broad electron beam of (c) moderate amplitude to (a) mitigate self-focusing, (b) avoid creating a wake in the blowout regime and (c) to create a 1D-like wake profile near the axis. Future work includes understanding frequency shift in the blowout regime, other strategies to mitigate self-focusing, and finding expressions for 3D modifications to the wake profile.

- **Optimization of electron beam, laser beam, and plasma density profile**

The profiles of the electron beam, laser beam, and tapered plasma density can all be tuned for greater frequency shift, better energy transfer efficiency, beam loading [112, 113], or other figures of merit.

- **Laser-driven photon acceleration**

A similar model to the PA ^{∞} differential equations (3.76-3.77) for a laser driver can be derived. While arbitrary frequency shifts may not be possible, significant frequency shifting of one pulse by another should.

- **investigating other methods of maintaining the driver and wake for extended photon acceleration**

In particular, combining a spatiotemporally shaped light or electron beams [90, 91, 111] with the PA^∞ model and density profile introduced in this dissertation could be useful in maintaining the PA^∞ driver over the distances needed for significant frequency shift. Spatiotemporally shaped light has been demonstrated in theory and simulations of electron [89] and photon [21] acceleration, and spatiotemporally matter has been demonstrated in theory and simulations of electron acceleration [111].

To evaluate the prospects of the unlimited photon acceleration scheme for attaining greater frequency shifts, consider the relation for frequency gain as a function of plasma density given in equation (3.85). In figure 4.1(a) are the plots of several curves indicating how the frequency shift as the laser propagates through decreasing plasma density are given for drive beams of charge $Q = 1, 2, \text{ and } 5$ nC, corresponding to amplitudes $A = 0.18, 0.35, \text{ and } 0.88$. In associating charge Q with amplitude A , the electron beam is assumed to have a root-mean-square radial extent $\sigma_r = 16 \mu\text{m}$. This is smaller than the radial extent $\sigma_r = 32 \mu\text{m}$ used in the quasi-3D simulation, chosen to match the laser beam waist $w_0 = 16 \mu\text{m}$. It is conceivable that a similarly sized drive beam could accelerate and guide a similarly sized laser pulse, or that the laser pulse could be focused down to a beam waist of $w_0 = 8 \mu\text{m}$. From the figure it can be seen that the stronger driver leads to greater frequency shifts for the same change in plasma density. Several points are marked on each trajectory indicating how far the laser and drive beam must propagate to reach that frequency-density point on the trajectory, if the initial laser pulse has wavelength $\lambda = 800 \mu\text{m}$ and the initial plasma density is $n_0 = 1.7 \times 10^{19} \text{ cm}^{-3}$. For a drive beam with $Q = 5$ nC, the model predicts $100\times$ frequency shift in 2 m of propagation. On the other hand, for a drive beam with $Q = 1$ nC a frequency shift of $100\times$ would require 1 km of propagation. The enhanced shifts with higher drive beam amplitude and more nonlinear wakes show the need for the nonlinear model derived in this chapter and the need to pursue higher-amplitude drivers in

subsequent research.

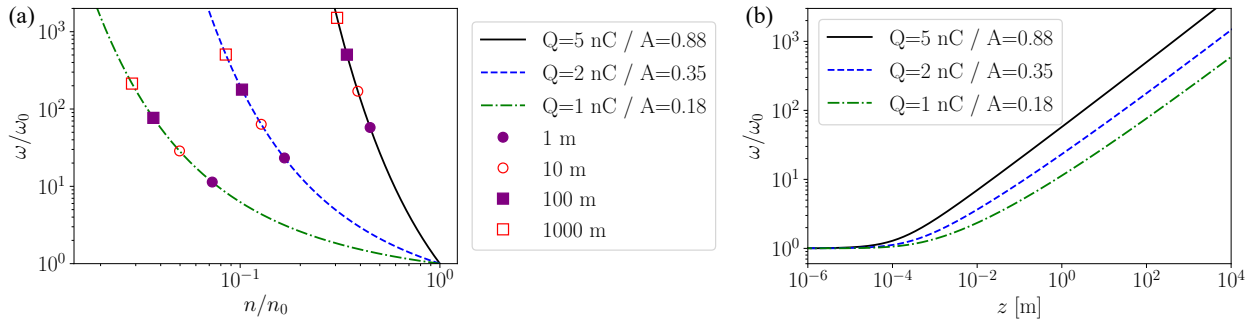


Figure 4.1: (a) Relative frequency gain as a function of density, $(\omega/\omega_0)(n)$, is plotted for beam drive charges $Q = 1$ nC, 2 nC, and 5 nC. The frequency gain-density curves are labeled with markers indicating the distance that the laser and drive beam must propagate to achieve the indicated shift. (b) Relative frequency gain as a function of z in meters, plotted for beam drive charges $Q = 1$ nC, 2 nC, and 5 nC.

In Fig. 4.1(b), the equation for ω as a function of n , Eqn. (3.85), is combined with the equation for z as a function of n , Eqn. (3.96), to plot relative frequency gain, ω/ω_0 , as a function of propagation distance, z . For a beam charge of $Q = 5$ nC, $100\times$ frequency shift is possible in less than 3 meters of propagation, $500\times$ frequency shift is possible in about 100 m, and $1000\times$ shift is possible in about 400 m. Beginning with the 2nd harmonic of a Ti:Sapph laser at 400 nm, a 1000 times shift (scaling the plasma / initial conditions to maintain similarity and fixing a_0) would result in pulses of coherent 4 Å (ångstrom) light. The photons in these pulses would have about 3 keV of energy each, well into the X-ray region of the electromagnetic spectrum. The pulse itself would have several joules of energy and would have a sub-femtosecond pulse duration if the energy enhancement and compression evidenced in this dissertation hold. To place these numbers in context, consider the ångstrom-wavelength free electron laser (FEL) at the Stanford Linear Accelerator (SLAC) facility [119], which produces X-ray pulses with 10^{12} photons per pulse, having wavelengths of 0.15 nm with energy $U = 3$ mJ and pulse duration $\tau = 100$ fs. The length of the SLAC FEL is about 1 km — the PA^∞ propagation length would be on the same order, about 400

m, but the PA^∞ X-ray pulse is $1000\times$ shorter and has $1000\times$ more energy.

The comparison becomes even more favorable if the pulse compression evidenced in the PA^∞ frequency shift is leveraged to stage acceleration. Noting that the frequency shift scales roughly as $\mathcal{O}(\sqrt{z})$ and hence that most shift occurs early in the propagation, more frequency shift can be attained in less propagation distance if the plasma density n is reset after a short time. Using a sawtooth density profile and restarting the PA^∞ scheme, it would be possible to achieve $100\times$ frequency shift in 3 meters, then $100\times$ again in 3 more meters, for $10000\times$ frequency shift in 10 m or less. In a smaller distance and at the same wavelength as the SLAC FEL, the staged unlimited photon acceleration scheme PA^∞ could generate X-ray beams $1000\times$ shorter and with $1000\times$ more photons and energy.

APPENDICES

APPENDIX A

On-line materials, including source code and animations

The FARSIGHT Vlasov-Poisson source code and documentation can be found on-line at <https://github.com/RTSandberg/FARSIGHT>.

In chapter II, several sequences of phase-space images were presented as part of the FARSIGHT test case studies. Each test case presented in this dissertation has an animation that can be found on YouTube:

- Weak Landau damping: <https://youtu.be/TTUCK9DrS1o>
- Strong Landau damping showing panels: <https://youtu.be/RH131FfbLms>
- Strong Landau damping: <https://youtu.be/lU-ed4AYQrM>
- Halo formation: <https://youtu.be/UIHV1ezdnFY>
- Warm two-stream instability: <https://youtu.be/vMXde63Nrec>
- Cold two-stream instability: <https://www.youtube.com/watch?v=rD-8xj-KJME>

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] James A. Rossmannith and David C. Seal. A positivity-preserving high-order semi-Lagrangian discontinuous Galerkin scheme for the Vlasov-Poisson equations. *J. Comput. Phys.*, 230(16):6203–6232, 2011.
- [2] David Sirajuddin and William N.G. Hitchon. A truly forward semi-Lagrangian WENO scheme for the Vlasov-Poisson system. *J. Comput. Phys.*, 2019.
- [3] Steven M. Lund, Alex Friedman, and Guillaume Bazouin. Sheet beam model for intense space charge: Application to Debye screening and the distribution of particle oscillation frequencies in a thermal equilibrium beam. *Physical Review Special Topics - Accelerators and Beams*, 14(5):054201, 5 2011.
- [4] Martin Campos Pinto, Eric Sonnendrücker, Alex Friedman, David P. Grote, and Steve M. Lund. Noiseless Vlasov-Poisson simulations with linearly transformed particles. *J. Comput. Phys.*, 275:236–256, 10 2014.
- [5] E. Esarey, C. B. Schroeder, and W. P. Leemans. Physics of laser-driven plasma-based electron accelerators. *Reviews of Modern Physics*, 81(3), 8 2009.
- [6] C. B. Schroeder, E. Esarey, C. G. R. Geddes, Cs. Toth, W. P. Leemans, Carl B. Schroeder, Wim Leemans, and Eric Esarey. Design considerations for a laser-plasma linear collider. In *AIP Conference Proceedings*. AIP, 2009.
- [7] S. Kneip, C. McGuffey, F. Dollar, M. S. Bloom, V. Chvykov, G. Kalintchenko, K. Krushelnick, A. Maksimchuk, S. P. D. Mangles, T. Matsuoka, Z. Najmudin, C. A. J. Palmer, J. Schreiber, W. Schumaker, A. G. R. Thomas, and V. Yanovsky. X-ray phase contrast imaging of biological specimens with femtosecond pulses of betatron radiation from a compact laser plasma wakefield accelerator. *Applied Physics Letters*, 99(9), 8 2011.
- [8] S. C. Wilks, J. M. Dawson, W. B. Mori, T. Katsouleas, and M. E. Jones. Photon accelerator. *Physical Review Letters*, 62(22), 5 1989.
- [9] A.A. Vlasov. The vibrational properties of an electron gas. *Journal of Experimental and Theoretical Physics*, 93(11):444–470, 11 1967.
- [10] Jose Tito Mendonca. Vlasov equation for photons and quasi-particles in a plasma. *The European Physical Journal D*, 68(4), 4 2014.

- [11] A. M. Hansen, K. L. Nguyen, B. J. Turnbull, D. Albright, R. K. Follett, R. Huff, J. Katz, D. Mastrosimone, A. L. Milder, L. Yin, J. P. Palastro, and D. H. Froula. Cross-Beam Energy Transfer Saturation by Ion Heating. *Physical Review Letters*, 126(7), 2 2021.
- [12] B. Svedung Wettervik, T. C. Dubois, and T. Fülöp. Vlasov modelling of laser-driven collisionless shock acceleration of protons. *Physics of Plasmas*, 23(5), 5 2016.
- [13] A. Grassi, L. Fedeli, A. Sgattoni, and A. Macchi. Vlasov simulation of laser-driven shock acceleration and ion turbulence. *Plasma Physics and Controlled Fusion*, 58(3), 2 2016.
- [14] I. H. Hutchinson. Electron holes in phase space: What they are and why they matter. *Physics of Plasmas*, 24(5), 5 2017.
- [15] W.B. Mori. The physics of the nonlinear optics of plasmas at relativistic intensities for short-pulse lasers. *IEEE Journal of Quantum Electronics*, 33(11):1942–1953, 1997.
- [16] Jose Tito Mendonca. *Theory of photon acceleration*. Institute of Physics Pub, 2001.
- [17] E. Esarey, A. Ting, and P. Sprangle. Frequency shifts induced in laser pulses by plasma waves. *Physical Review A*, 42(6):3526–3531, 9 1990.
- [18] V. A. Mironov, A. M. Sergeev, E. V. Vanin, G. Brodin, and J. Lundberg. Upper limits for frequency up-conversion in the nonlinear photon accelerator. *Physical Review A*, 46(10):R6178–R6180, 11 1992.
- [19] C. D. Murphy, R. Trines, J. Vieira, A. J.W. Reitsma, R. Bingham, J. L. Collier, E. J. Divall, P. S. Foster, C. J. Hooker, A. J. Langley, P. A. Norreys, R. A. Fonseca, F. Fiuza, L. O. Silva, J. T. Mendonça, W. B. Mori, J. G. Gallacher, R. Viskup, D. A. Jaroszynski, S. P.D. Mangles, A. G.R. Thomas, K. Krushelnick, and Z. Najmudin. Evidence of photon acceleration by laser wake fields. *Physics of Plasmas*, 13(3), 2006.
- [20] Matthew R. Edwards, Kenan Qu, Qing Jia, Julia M. Mikhailova, and Nathaniel J. Fisch. Cascaded chirped photon acceleration for efficient frequency conversion. *Physics of Plasmas*, 25(5), 5 2018.
- [21] A. J. Howard, D. Turnbull, A. S. Davies, P. Franke, D. H. Froula, and J. P. Palastro. Photon Acceleration in a Flying Focus. *Physical Review Letters*, 123(12):124801, 9 2019.
- [22] Zan Nie, Chih Hao Pai, Jianfei Hua, Chaojie Zhang, Yipeng Wu, Yang Wan, Fei Li, Jie Zhang, Zhi Cheng, Qianqian Su, Shuang Liu, Yue Ma, Xiaonan Ning, Yunxiao He, Wei Lu, Hsu Hsin Chu, Jyhpyng Wang, Warren B. Mori, and Chan Joshi. Relativistic single-cycle tunable infrared pulses generated from a tailored plasma density structure. *Nature Photonics*, 12(8):489–494, 8 2018.
- [23] Eric Sonnendrucker, John J Barnard, Alex Friedman, David P Grote, and Steve M Lund. Simulation of heavy ion beams with a semi-Lagrangian Vlasov solver. *Nuclear Instruments and Methods in Physics Research A*, 464:470–476, 2001.

- [24] Bedros Afeyan, Fernando Casas, Nicolas Crouseilles, Adila Dodhy, Erwan Faou, Michel Mehrenberger, and Eric Sonnendrücker. Simulations of kinetic electrostatic electron nonlinear (KEEN) waves with variable velocity resolution grids and high-order time-splitting. *The European Physical Journal D*, 68(10):295, 10 2014.
- [25] B. A. Shadwick, G. M. Tarkenton, E. Esarey, and C. B. Schroeder. Fluid and Vlasov models of low-temperature, collisionless, relativistic plasma interactions. *Physics of Plasmas*, 12(5):056710, 5 2005.
- [26] A. G.R. Thomas. Vlasov simulations of thermal plasma waves with relativistic phase velocity in a Lorentz boosted frame. *Physical Review E*, 94(5), 11 2016.
- [27] Charles K Birdsall and A Bruce Langdon. *Plasma Physics via Computer Simulation*. IOP Publishing, 1991.
- [28] R Hockney and J Eastwood. *Computer Simulation Using Particles*. Taylor & Francis, 1988.
- [29] W. M. Nevins, G. W. Hammett, A. M. Dimits, W. Dorland, and D. E. Shumaker. Discrete particle noise in particle-in-cell simulations of plasma microturbulence. *Physics of Plasmas*, 12(12):1–16, 2005.
- [30] Dennis W. Hewett. Fragmentation, merging, and internal dynamics for PIC simulation with finite size particles. *J. Comput. Phys.*, 189(2):390–426, 8 2003.
- [31] David J. Larson and Christopher V. Young. A finite mass based method for Vlasov-Poisson simulations. *J. Comput. Phys.*, 284:171–185, 3 2015.
- [32] A.J. Christlieb, R. Krasny, J.P. Verboncoeur, J.W. Emhoff, and I.D. Boyd. Grid-free plasma simulation techniques. *IEEE Transactions on Plasma Science*, 34(2):149–165, 4 2006.
- [33] D M Thomas and J T Holgate. A treecode to simulate dust–plasma interactions. *Plasma Physics and Controlled Fusion*, 59(2):025002, dec 2016.
- [34] Torsten Kessler, Sergej Rjasanow, and Steffen Weisser. Vlasov-Poisson system tackled by particle simulation utilizing boundary element methods. *SIAM Journal on Scientific Computing*, 42(1):B299–B326, 2020.
- [35] Pavel I. Shustov, Ilya V. Kuzichev, Ivan Y. Vasko, Anton V. Artemyev, and Andrew J. Gerrard. The dynamics of electron holes in current sheets. *Physics of Plasmas*, 28(1), 1 2021.
- [36] Martin V. Goldman, David L. Newman, and Philip Pritchett. Vlasov simulations of electron holes driven by particle distributions from PIC reconnection simulations with a guide field. *Geophysical Research Letters*, 35(22), 11 2008.
- [37] Jeffrey William Banks and Jeffrey Alan Furst Hittinger. A new class of nonlinear finite-volume methods for Vlasov simulation. *IEEE Transactions on Plasma Science*, 38(9 PART 1):2198–2207, 9 2010.

- [38] G.V. Vogman, U. Shumlak, and P. Colella. Conservative fourth-order finite-volume Vlasov–Poisson solver for axisymmetric plasmas in cylindrical (r,v,v) phase space coordinates. *Journal of Computational Physics*, 373, 11 2018.
- [39] R.E. Heath, I.M. Gamba, P.J. Morrison, and C. Michler. A discontinuous Galerkin method for the Vlasov–Poisson system. *Journal of Computational Physics*, 231(4), 2 2012.
- [40] Joseph T. Parker and Paul J. Dellar. Fourier–Hermite spectral representation for the Vlasov–Poisson system in the weakly collisional limit. *Journal of Plasma Physics*, 81(2), 4 2015.
- [41] E. Camporeale, G.L. Delzanno, B.K. Bergen, and J.D. Moulton. On the velocity space discretization for the Vlasov–Poisson system: Comparison between implicit Hermite spectral and Particle-in-Cell methods. *Computer Physics Communications*, 198, 1 2016.
- [42] J. Denavit. Numerical simulation of plasmas with periodic smoothing in phase space. *J. Comput. Phys.*, 9(1):75–98, 1972.
- [43] C Z Cheng and Georg Knorr. The Integration of the Vlasov Equation in Configuration Space. *J. Comput. Phys.*, 22:330–351, 1976.
- [44] Eric Sonnendrücker, Jean Roche, Pierre Bertrand, and Alain Ghizzo. The Semi-Lagrangian Method for the Numerical Resolution of the Vlasov Equation. *J. Comput. Phys.*, 149:201–220, 1999.
- [45] Jing Mei Qiu and Chi Wang Shu. Positivity preserving semi-Lagrangian discontinuous Galerkin formulation: Theoretical analysis and application to the Vlasov-Poisson system. *J. Comput. Phys.*, 230(23):8386–8409, 2011.
- [46] Yingda Cheng, Irene M Gamba, and Philip J Morrison. Study of conservation and recurrence of Runge-Kutta discontinuous Galerkin schemes for Vlasov-Poisson systems. *Journal of Scientific Computing*, 56:319–349, 2013.
- [47] Chieh Sen Huang, Todd Arbogast, and Chen Hui Hung. A semi-Lagrangian finite difference WENO scheme for scalar nonlinear conservation laws. *J. Comput. Phys.*, 322:559–585, 10 2016.
- [48] Andrew Christlieb, Matthew Link, Hyoseon Yang, and Ruimeng Chang. High-Order Semi-Lagrangian WENO Schemes Based on Non-polynomial Space for the Vlasov Equation. *Commun. Appl. Math. Comput.*, 2021.
- [49] Xiaofeng Cai, Wei Guo, and Jing Mei Qiu. A high order semi-Lagrangian discontinuous Galerkin method for Vlasov–Poisson simulations without operator splitting. *J. Comput. Phys.*, 354:529–551, 2 2018.

- [50] Tao Xiong, Giovanni Russo, and Jing Mei Qiu. Conservative Multi-dimensional Semi-Lagrangian Finite Difference Scheme: Stability and Applications to the Kinetic and Fluid Simulations. *Journal of Scientific Computing*, 79(2):1241–1270, 5 2019.
- [51] C. Alard and S. Colombi. A cloudy Vlasov solution. *Monthly Notices of the Royal Astronomical Society*, 359(1):123–163, 5 2005.
- [52] Stéphane Colombi and Christophe Alard. A ‘metric’ semi-Lagrangian Vlasov–Poisson solver. *Journal of Plasma Physics*, 83(3), 6 2017.
- [53] B. Wang, G. H. Miller, and P. Colella. A particle-in-cell method with adaptive phase-space remapping for kinetic plasmas. *SIAM Journal on Scientific Computing*, 33(6):3509–3537, 2011.
- [54] J.A.F. Hittinger and J.W. Banks. Block-structured adaptive mesh refinement algorithms for Vlasov simulation. *Journal of Computational Physics*, 241, 5 2013.
- [55] M. Gutnic, M. Haefele, I. Paun, and E. Sonnendrücker. Vlasov simulations on an adaptive phase-space grid. *Computer Physics Communications*, 164:214–219, 12 2004.
- [56] N. Besse, E. Deriaz, and Madaule. Adaptive multiresolution semi-Lagrangian discontinuous Galerkin methods for the Vlasov equations. *J. Comput. Phys.*, 332:376–417, 3 2017.
- [57] Wei Guo and Yingda Cheng. An adaptive multiresolution Discontinuous Galerkin method for time-dependent transport equations in multidimensions. *SIAM J. Sci. Comput.*, 39(6):A2962–A2992, 2017.
- [58] Nathan Vaughn, Leighton Wilson, and Robert Krasny. A GPU-Accelerated Barycentric Lagrange Treecode. In *21st IEEE Int. Workshop Parallel Distrib. Sci. Eng. Comput. (PDSEC)*, pages 701–710, 2020.
- [59] Peter Bosler, Lei Wang, Christiane Jablonowski, and Robert Krasny. A Lagrangian particle/panel method for the barotropic vorticity equations on a rotating sphere. *Fluid Dynamics Research*, 46(3), 6 2014.
- [60] Alexandre Joel Chorin. Numerical study of slightly viscous flow. *J. Fluid Mech.*, 57(4):785–796, 3 1973.
- [61] Robert Krasny. Desingularization of periodic vortex sheet roll-up. *J. Comput. Phys.*, 65:292–313, 1986.
- [62] G Cottet and P Koumoutsakos. *Vortex Methods: Theory and Practice*. Cambridge University Press, 2000.
- [63] Mirta Perlman. On the accuracy of vortex methods. *J. Comput. Phys.*, 59(2):200–223, 6 1985.
- [64] P. Koumoutsakos. Inviscid axisymmetrization of an elliptical vortex. *J. Comput. Phys.*, 138:821–857, 1997.

- [65] A. Myers, P. Colella, and B. Van Straalen. A 4th-order particle-in-cell method with phase-space remapping for the Vlasov-Poisson equation. *SIAM J. Sci. Comput.*, 39:B467–B485, 2017.
- [66] Josh Barnes and Piet Hut. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature*, 324(4):446–449, 1986.
- [67] Lei Wang, Robert Krasny, and Svetlana Tlupova. A kernel-independent treecode based on barycentric Lagrange interpolation. *Commun. Comput. Phys.*, 28:1415–1436, 2020.
- [68] J.-P. Berrut and L. N. Trefethen. Barycentric Lagrange interpolation. *SIAM Rev.*, 46(3):501–517, 2004.
- [69] Lev Landau. On the vibrations of the electron plasma. volume 10, pages 25–34. 1946.
- [70] C. Villani. Particle systems and nonlinear Landau damping. *Phys. Plasmas*, 21:030901, 2014.
- [71] J. Thomas Beale and Andrew Majda. Vortex Methods. II: Higher Order Accuracy in Two and Three Dimensions. *Mathematics of Computation*, 39(159), 7 1982.
- [72] P. A. Raviart. An analysis of particle methods. In Franco Brezzi, editor, *Numerical Methods in Fluid Dynamics*, pages 243–324. Springer Berlin Heidelberg, Berlin, 1985.
- [73] Nicolas Crouseilles, Michel Mehrenberger, and Eric Sonnendrücker. Conservative semi-Lagrangian schemes for Vlasov equations. *J. Comput. Phys.*, 229(6):1927–1953, 3 2010.
- [74] Jing Mei Qiu and Chi Wang Shu. Conservative semi-Lagrangian finite difference WENO formulations with applications to the vlasov equation. *Communications in Computational Physics*, 2011.
- [75] Xiaofeng Cai, Jianxian Qiu, and Jing Mei Qiu. A conservative semi-Lagrangian HWENO method for the Vlasov equation. *J. Comput. Phys.*, 323:95–114, 10 2016.
- [76] Xiaofeng Cai, Sebastiano Boscarino, and Jing-Mei Qiu. High order semi-lagrangian discontinuous galerkin method coupled with runge-kutta exponential integrators for nonlinear vlasov dynamics. *Journal of Computational Physics*, 427:110036, 2021.
- [77] Leighton Wilson, Nathan Vaughn, and Robert Krasny. A GPU-accelerated fast multipole method based on barycentric Lagrange interpolation and dual tree traversal. *Comput. Phys. Commun.*, 265:108017, 2021.
- [78] Jörg Rossbach, Jochen R. Schneider, and Wilfried Wurth. 10 years of pioneering x-ray science at the free-electron laser flash at desy. *Physics Reports*, 808:1–74, 2019. 10 years of pioneering X-ray science at the Free-Electron Laser FLASH at DESY.
- [79] M A Khokhlova and V V Strelkov. Highly efficient XUV generation via high-order frequency mixing. *New Journal of Physics*, 22(9):093030, sep 2020.

- [80] Zenghu Chang, Andy Rundquist, Haiwen Wang, Margaret M. Murnane, and Henry C. Kapteyn. Generation of coherent soft x rays at 2.7 nm using high harmonics. *Phys. Rev. Lett.*, 79:2967–2970, Oct 1997.
- [81] B. R. Benware, C. D. Macchietto, C. H. Moreno, and J. J. Rocca. Demonstration of a high average power tabletop soft x-ray laser. *Phys. Rev. Lett.*, 81:5804–5807, Dec 1998.
- [82] T. Esirkepov, S. V. Bulanov, M. Yamagiwa, and T. Tajima. Electron, positron, and photon wakefield acceleration: Trapping, wake overtaking, and ponderomotive acceleration. *Physical Review Letters*, 96(1):014803, 1 2006.
- [83] E. Esarey, G. Joyce, and P. Sprangle. Frequency up-shifting of laser pulses by copropagating ionization fronts. *Physical Review A*, 44(6):3908–3911, 9 1991.
- [84] Luis Oliveira e Silva and Jose Tito Mendonca. Photon acceleration in superluminous and accelerated ionization fronts. *IEEE Transactions on Plasma Science*, 24(2):316–322, 4 1996.
- [85] Maxim R. Shcherbakov, Kevin Werner, Zhiyuan Fan, Noah Talisa, Enam Chowdhury, and Gennady Shvets. Photon acceleration and tunable broadband harmonics generation in nonlinear time-dependent metasurfaces. *Nature Communications*, 10(1):1345, 2019.
- [86] J. M. Dias, C. Stenz, N. Lopes, X. Badiche, F. Blasco, A. Dos Santos, L. Oliveira e Silva, A. Mysyrowicz, A. Antonetti, and J. T. Mendonça. Experimental evidence of photon acceleration of ultrashort laser pulses in relativistic ionization fronts. *Phys. Rev. Lett.*, 78:4773–4776, Jun 1997.
- [87] C. W. Siders, S. P. Le Blanc, D. Fisher, T. Tajima, M. C. Downer, A. Babine, A. Stepanov, and A. Sergeev. Laser wakefield excitation and measurement by femtosecond longitudinal interferometry. *Phys. Rev. Lett.*, 76:3570–3573, May 1996.
- [88] C. D. Murphy, R. Trines, J. Vieira, A. J. W. Reitsma, R. Bingham, J. L. Collier, E. J. Divall, P. S. Foster, C. J. Hooker, A. J. Langley, P. A. Norreys, R. A. Fonseca, F. Fiuza, L. O. Silva, J. T. Mendonça, W. B. Mori, J. G. Gallacher, R. Viskup, D. A. Jaroszynski, S. P. D. Mangles, A. G. R. Thomas, K. Krushelnick, and Z. Najmudin. Evidence of photon acceleration by laser wake fields. *Physics of Plasmas*, 13(3):033108, 2006.
- [89] J. P. Palastro, J. L. Shaw, P. Franke, D. Ramsey, T. T. Simpson, and D. H. Froula. Dephasingless laser wakefield acceleration. *Phys. Rev. Lett.*, 124:134802, Mar 2020.
- [90] Alexander Debus, Richard Pausch, Axel Huebl, Klaus Steiniger, René Widera, Thomas E. Cowan, Ulrich Schramm, and Michael Bussmann. Circumventing the Dephasing and Depletion Limits of Laser-Wakefield Acceleration. *Physical Review X*, 9(3), 9 2019.

- [91] C. Caizergues, S. Smartsev, V. Malka, and C. Thaury. Phase-locked laser-wakefield electron acceleration. *Nature Photonics*, 14(8):475–479, 2020.
- [92] S. V. Bulanov, V. A. Vshivkov, G. I. Dudnikova, N. M. Naumova, F. Pegoraro, and I. V. Pogorelsky. Laser acceleration of charged particles in inhomogeneous plasmas i. 23(4):259–269, 1997.
- [93] P. Sprangle, B. Hafizi, J. R. Peñano, R. F. Hubbard, A. Ting, C. I. Moore, D. F. Gordon, A. Zigler, D. Kaganovich, and T. M. Antonsen. Wakefield generation and GeV acceleration in tapered plasma channels. *Physical Review E*, 63(5):056405, 4 2001.
- [94] E. Guillaume, A. Döpp, C. Thaury, K. Ta Phuoc, A. Lifschitz, G. Grittani, J.-P. Goddet, A. Tafzi, S. W. Chou, L. Veisz, and V. Malka. Electron Rephasing in a Laser-Wakefield Accelerator. *Physical Review Letters*, 115(15):155002, 10 2015.
- [95] S. V. Bulanov, T. Zh. Esirkepov, Y. Hayashi, H. Kiriyama, J. K. Koga, H. Kotaki, M. Mori, and M. Kando. On some theoretical problems of laser wake-field accelerators. *Journal of Plasma Physics*, 82(3):905820308, 6 2016.
- [96] Steven Weinberg. Eikonal Method in Magnetohydrodynamics. *Physical Review*, 126(6):1899–1909, 6 1962.
- [97] D. Guénot, D. Gustas, A. Vernier, B. Beaurepaire, F. Böhle, M. Bocoum, M. Lozano, A. Jullien, R. Lopez-Martens, A. Lifschitz, and J. Faure. Relativistic electron beams driven by kHz single-cycle light pulses. *Nature Photonics*, 11(5):293–296, 5 2017.
- [98] J. Kim, V. L. J. Phung, K. Roh, M. Kim, K. Kang, and H. Suk. Development of a density-tapered capillary gas cell for laser wakefield acceleration. *Review of Scientific Instruments*, 92(2):023511, 2 2021.
- [99] A.I. Akhiezer and R.V. Polovin. Theory of wave motion of an electron plasma. *Journal of Experimental and Theoretical Physics*, 3(5):696–705, 1956.
- [100] E. Infeld and G. Rowlands. Relativistic bursts. *Physical Review Letters*, 62(10), 3 1989.
- [101] Prabal Singh Verma, Sudip Sengupta, and Predhiman Kaw. Bernstein-Greene-Kruskal waves in relativistic cold plasma. *Physics of Plasmas*, 19(3), 3 2012.
- [102] Prabal Singh Verma, Sudip Sengupta, and Predhiman Kaw. Breaking of longitudinal Akhiezer-Polovin waves. *Physical Review Letters*, 108(12), 3 2012.
- [103] J. B. Rosenzweig. Nonlinear plasma dynamics in the plasma wake-field accelerator. *Physical Review Letters*, 58(6):555–558, 2 1987.
- [104] Paul F. Byrd and Morris D. Friedman. *Handbook of Elliptic Integrals for Engineers and Scientists*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1971.

- [105] Ivar Bendixson. Sur les courbes définies par des équations différentielles. *Acta Mathematica*, 24(0):1–88, 1901.
- [106] W. Lu, C. Huang, M. Zhou, W. B. Mori, and T. Katsouleas. Nonlinear Theory for Relativistic Plasma Wakefields in the Blowout Regime. *Physical Review Letters*, 96(16):165002, 4 2006.
- [107] Roy G. Hemker. *Particle-in-Cell Modeling of Plasma-Based Accelerators in Two and Three Dimensions*. PhD thesis, UCLA, 1999.
- [108] R. A. Fonseca, L. O. Silva, F. S. Tsung, V. K. Decyk, W. Lu, C. Ren, W. B. Mori, S. Deng, S. Lee, T. Katsouleas, and J. C. Adam. Osiris: A three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators. In Peter M. A. Sloot, Alfons G. Hoekstra, C. J. Kenneth Tan, and Jack J. Dongarra, editors, *Computational Science — ICCS 2002*, pages 342–351, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [109] Rémi Lehe, Manuel Kirchen, Igor A. Andriyash, Brendan B. Godfrey, and Jean Luc Vay. A spectral, quasi-cylindrical and dispersion-free Particle-In-Cell algorithm. *Computer Physics Communications*, 203:66–82, 6 2016.
- [110] V. Yakimenko, L. Alsberg, E. Bong, G. Bouchard, C. Clarke, C. Emma, S. Green, C. Hast, M. J. Hogan, J. Seabury, N. Lipkowitz, B. O’Shea, D. Storey, G. White, and G. Yocky. FACET-II facility for advanced accelerator experimental tests. *Physical Review Accelerators and Beams*, 22(10):101301, 10 2019.
- [111] Fei Li, Thamine N. Dalichaouch, Jacob R. Pierce, Xinlu Xu, Frank S. Tsung, Wei Lu, Chan Joshi, and Warren B. Mori. Ultra-bright electron bunch injection in a plasma wakefield driven by a superluminal flying focus electron beam, 2021.
- [112] T. Katsouleas, Scott Wilks, Pisin Chen, J. M. Dawson, and J. J. Su. Beam loading in plasma accelerators. *Particle Accelerators*, 22:81–99, 9 1987.
- [113] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. Beam loading by electrons in nonlinear plasma wakes. *Physics of Plasmas*, 16(5):056705, 5 2009.
- [114] I. Y. Dodin and N. J. Fisch. On the nature of kinetic electrostatic electron nonlinear (KEEN) waves. *Physics of Plasmas*, 21(3):034501, 3 2014.
- [115] Yingda Cheng, Irene M. Gamba, and Philip J. Morrison. Study of conservation and recurrence of Runge-Kutta discontinuous Galerkin schemes for Vlasov-Poisson systems. *Journal of Scientific Computing*, 56(2):319–349, 8 2013.
- [116] T. W. Johnston, Y. Tyshetskiy, A. Ghizzo, and P. Bertrand. Persistent subplasma-frequency kinetic electrostatic electron nonlinear waves. *Physics of Plasmas*, 16(4):042105, 4 2009.

- [117] Roman V. Shcherbakov. Dispersion of waves in relativistic plasmas with isotropic particle distributions. *Physics of Plasmas*, 16(3):032104, 3 2009.
- [118] E. W. Laing and D. A. Diver. Relativistic Landau damping of longitudinal waves in isotropic pair plasmas. *Physics of Plasmas*, 13(9):092115, 9 2006.
- [119] P. Emma, R. Akre, J. Arthur, R. Bionta, C. Bostedt, J. Bozek, A. Brachmann, P. Bucksbaum, R. Coffee, F.-J. Decker, Y. Ding, D. Dowell, S. Edstrom, A. Fisher, J. Frisch, S. Gilevich, J. Hastings, G. Hays, Ph. Hering, Z. Huang, R. Iverson, H. Loos, M. Messerschmidt, A. Miahnahri, S. Moeller, H.-D. Nuhn, G. Pile, D. Ratner, J. Rzepiela, D. Schultz, T. Smith, P. Stefan, H. Tompkins, J. Turner, J. Welch, W. White, J. Wu, G. Yocky, and J. Galayda. First lasing and operation of an ångstrom-wavelength free-electron laser. *Nature Photonics*, 4(9):641–647, 9 2010.