

AFRL-AFOSR-VA-TR-2023-0348

2018 - Modeling of Ultrafast High-Current Electron Beam Transport in Air

Thomas, Alexander
REGENTS OF THE UNIVERSITY OF MICHIGAN
503 THOMPSON STREET
ANN ARBOR, MI, 48109
USA

04/25/2023
Final Technical Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
20230425	Final	START DATE 20190201	END DATE 20230131

4. TITLE AND SUBTITLE
2018 - Modeling of Ultrafast High-Current Electron Beam Transport in Air

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
	FA9550-19-1-0072	61102F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Alexander Thomas

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
REGENTS OF THE UNIVERSITY OF MICHIGAN
503 THOMPSON STREET
ANN ARBOR, MI 48109
USA

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
Air Force Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203	AFRL/AFOSR RTB1	AFRL-AFOSR-VA-TR-2023-0348

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The overall goal of this project was to understand the physics of the generation and propagation of ultrashort electron beams in atmosphere using advanced simulation tools relevant to experiments on the PHEENIX laser at AFRL Kirtland. Propagation of electron beams in air has been of interest to the Department of Defense for decades. The use of electron beams to disrupt shielded electronic components at a distance is a possible application, but a large body of prior research shows that beams are susceptible to instabilities that prevent it being deployed effectively. However, the ultrashort duration of beams generated by plasma wakefield accelerators means that the bunch duration is much shorter than the Langmuir period in atmospheric density plasma, which means that some instabilities may be mitigated for beam propagation because they are mediated at the plasma frequency. The ability to easily generate beams in excess of 1 GeV may enable penetration of radiation hardened targets. With a highly compact set-up and pencil beam potentially completely directional over all angles, LWFA allows more flexibility compared with a conventional accelerator. However, in additional to addressing technical challenges associated with the source, such as the repetition rate, there is limited experimental data on the propagation of such femtosecond duration beams in plasma or air. In work proposed here, we worked in collaboration with the AFRL group on computational model development, physics and simulation of the beam propagation. As outcomes from the project, we have implemented a new algorithm for simulating collisional ionization within a particle-in-cell framework. Our method treats and calculates ionization events deterministically as each particle's contribution to the ionization rate is calculated explicitly and deposited onto a grid. This algorithm is now fully implemented within the PIC code framework OSIRIS 4.0, and we performed several simulations high-lighting the new physics that is captured when considering collisional ionization in different scenarios. These include collimated proton beams generated in laser interactions with water targets and simulations of experimental measurements of the evolution of electric field from optical-field ionized plasmas using electron bunches produced from a high repetition-rate laser plasma accelerator. We have also developed 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 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 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 GPUaccelerated barycentric Lagrange treecode.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	c. THIS PAGE		
U	U	U	UU	76

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
JOHN LUGINSLAND	000-0000

Standard Form 298 (Rev.5/2020)
Prescribed by ANSI Std. Z39.18

DISTRIBUTION A: Distribution approved for public release.

To: technicalreports@afosr.af.mil
Subject: Final Progress Report

Contract/Grant Title: Modeling of Ultrafast High-Current Electron Beam Transport in Air
Contract/Grant #: FA9550-19-1-0072
Reporting Period: 2/1/19–1/31/2023
Principal Investigator: Alexander Thomas

Final Report

Accomplishments

Research Objectives

The overall goal of this project was to understand the physics of the generation and propagation of ultrashort electron beams in atmosphere using advanced simulation tools relevant to experiments on the PHEENIX laser at AFRL Kirtland. Propagation of electron beams in air has been of interest to the Department of Defense for decades. The use of electron beams to disrupt shielded electronic components at a distance is a possible application, but a large body of prior research shows that beams are susceptible to instabilities that prevent it being deployed effectively. However, the ultrashort duration of beams generated by plasma wakefield accelerators means that the bunch duration is much shorter than the Langmuir period in atmospheric density plasma, which means that some instabilities may be mitigated for beam propagation because they are mediated at the plasma frequency. The ability to easily generate beams in excess of 1 GeV may enable penetration of radiation hardened targets. With a highly compact set-up and pencil beam potentially completely directional over 4π, LWFA allows more flexibility compared with a conventional accelerator. However, in additional to addressing technical challenges associated with the source, such as the repetition rate, there is limited experimental data on the propagation of such femtosecond duration beams in plasma or air. In work proposed here, we worked in collaboration with the AFRL group on computational model development, physics and simulation of the beam propagation.

Our project had two main objectives:

DISTRIBUTION A: Distribution approved for public release.

- Investigate the physics of wakefield-accelerated femtosecond-duration beam propagation in atmospheric-density plasma / air through simulations using the particle-in-cell code OSIRIS 4.0. This will require some developments to the physics models to correctly model collisions and ionization. We will study instability development and potential mitigation and benchmark with experimental results in collaboration with the AFRL group.
- Develop mesh-free particle method for accurate plasma dynamics simulation based on a treecode-accelerated field solver and an adaptively refined particle/panel representation for the phase-space density. These should be able to augment existing particle-in-cell frameworks to give higher fidelity modeling of the fine-scale phase-space structure that arises in highly nonlinear plasma wakefields that is crucial to the long time dynamical development.

Major activities

We have implemented a new algorithm for simulating collisional ionization within a particle-in-cell framework. Specifically, we steer away from treating these ionization events as particle- particle interactions that occur via a random process but instead go back to transferring information between particles and grids. Our method treats and calculates ionization events deterministically as each particle's contribution to the ionization rate is calculated explicitly and deposited onto a grid. This grid of ionization rates is then used to advance ion densities, which allows us to track how much new charge is generated each timestep, so we can create newly ionized electrons accordingly. Additionally, the ionization rate grid, with little modification, keeps track of how much energy is lost per cell due to ionization physics. We interpolate this information back onto the particles; this allows for a continuous decrease in the energy of macro-particles as they participate in ionization events and allows for the easy calculation of the new momentum of ionized electrons. Collectively, these particle-to-grid and grid-to-particle information transfers act as a "smoothing" process, reducing noise considerably compared to other current algorithms.

This algorithm is now fully implemented within the PIC code OSIRIS, and we performed several simulations high- lighting the new physics that is captured when considering collisional ionization in different scenarios. These

DISTRIBUTION A: Distribution approved for public release.

include some ongoing work in collaboration with Charlotte Palmer (Queen's University Belfast, UK) on collimated proton beams generated in laser interactions with water targets, for which ionization effects play an important role and simulations of experimental measurements of the evolution of electric field from optical-field ionized plasmas using electron bunches produced from a high repetition-rate laser plasma accelerator. Papers from these studies are in preparation for publication.

We have also developed 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 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.

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; 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 $O(N^2)$ operations using direct summation, but the cost is reduced to $O(N \log N)$ 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.

DISTRIBUTION A: Distribution approved for public release.

Specific objectives

For the first objective of studying relativistic collisional ionization, we wrote and tested a new collisional ionization package for particle-in-cell codes, which was implemented in the OSIRIS 4.0 framework. This has the following features

1. It calculates rate equations for collisional ionization on a spatial grid, interpolating the particle contributions to the rate using the regular PIC weighting functions. This results in a smoother calculation of the rate equations for lower particle numbers.

2. The rates are calculated with semi-analytic relativistically correct cross sections for impact ionization [1] that cover a wide range of elements. The data for the binding energies was also used for calculating the field ionization rates in OSIRIS, meaning the field ionization model was extended to any element (data was previously only for a few particular low-Z elements).

3. Momentum conservation is calculated by storing the accumulated momentum density on a grid, which is transferred to ionized electrons when they are created. Moving ions for ionizable species was also implemented (it was not previously in the OSIRIS framework).

For the second objective of developing a mesh-free particle method for accurate plasma dynamics simulation based on a treecode-accelerated field solver and an adaptively refined particle/panel representation for the phase-space density, we have achieved

1. Implementing the barycentric Lagrange version of the treecode for the Poisson equation.

2. Formulated a regularized and periodic electric field kernel.

3. Added adaptively refined quadrilateral panel mesh refinement in phase space.

4. Ported the code to C++ with GPU acceleration.

DISTRIBUTION A: Distribution approved for public release.

Other achievements

An additional new study by student Ryan Sandberg was to simulate using the plasma wakefield to generate intense XUV radiation. The propagating density gradients of the wakefield may be used to frequency upshift a trailing witness laser pulse, a process known as 'photon acceleration'. In a uniform plasma, the witness laser will eventually dephase from the accelerating gradient because of group delay in the plasma. Here, we find the phase matching conditions for the pulse using a tapered density profile. An analytic solution for a 1d nonlinear plasma wake with an electron beam driver indicates that, even though the plasma density decreases, the frequency shift reaches no asymptotic limit, i.e. is unlimited provided the wake can be sustained. In quasi-3d particle-in-cell simulations using the calculate density profile we demonstrate frequency shifts of the witness laser beam up to ten times, limited only by simulation resolution and non-optimal driver evolution. In 1d simulations, more than 30 times shifts were observed. The pulse energy increases in this process, by a factor of five, the wake guides the pulse and it is temporally compressed by group velocity dispersion, resulting in the resulting XUV laser pulse having relativistic ($a_0 \sim 1$) intensity.

Dissemination of results to communities of interest

Published papers

- F. Albert et al., *2020 Roadmap on Plasma Accelerators* New J. Phys (2020) — A.G.R. Thomas co-editor and author of first section in collection.

- Russell, Brandon K.; Campbell, Paul T.; Thomas, Alexander G. R.; Willingale, Louise, *Multiple species laser-driven ion-shock acceleration*, Plasma Physics and Controlled Fusion, (2021).

- Ma, Y.; Seipt, D.; Krushelnick, K.; Thomas, A. G. R., *Generation of straight and curved hollow plasma channels by laser-generated nonlinear wakefields and studies of ultra-intense laser pulse guiding*, Phys. Plasmas 28, 063104 (2021).

- Ma, Y.; Seipt, D.; Hussein, A. E.; Hakimi, S.; Beier, N. F.; Hansen, S. B.; Hinojosa, J.; Maksimchuk, A.; Nees, J.; Krushelnick, K.; Thomas, A. G. R.; Dollar, F., *The effects of laser polarization and wavelength*

DISTRIBUTION A: Distribution approved for public release.

on injection dynamics of a laser wakefield accelerator, Phys. Plasmas 28, 063101 (2021).

- A.G.R. Thomas and D. Seipt, *Modelling Chromatic Emittance Growth in Staged Plasma Wakefield Acceleration to 1 TeV using Nonlinear Transfer Matrices*, Phys. Rev. Accel Beams 24, 104602, 2021.

- Ryan Sandberg and Alexander G.R. Thomas, *Photon Acceleration from Optical to XUV*, Physical Review Letters 130, 085001, 2023.

Manuscripts in preparation

- Ryan Sandberg and Alexander G.R. Thomas, *X-ray photon ultrashort pulses with matched phase* **follow up paper to Phys. Rev. Lett. in preparation**.

- Ryan Sandberg, Robert Krasny and Alexander G.R. Thomas, *FARSIGHT: A Forward Adaptively Refined and Regularized Semi-Lagrangian Integral Green's function Hierarchical Tree-code accelerated method for the Vlasov-Poisson system.* **In preparation for J. Comput. Phys.**

- Stephen DiIorio and Alexander G. R. Thomas, *A Collisional Ionization Physics Package for Particle-in-Cell Codes* **in preparation.**

- S. DiIorio, Z.-H. He, A. S. Joglekar, J. A. Nees, B. Hou, R. Clarke, K. Krushelnick, and A. G. R. Thomas, *Ultrafast Probing of Non-equilibrium Plasmas Using Laser-Wakefield-Accelerated Electron Bunches.* **in preparation.**

- C. A. J. Palmer1, M. J. V. Streeter1, B. Loughran1, H. Ahmed2, S. Astbury2, M. Balcazar5, M. Borghesi1, N. Bourgeois2, C. B. Curry3,4, S. J. D. Dann2, S. DiIorio5, N. P. Dover6, T. Dzelzainis2, O. Ettlinger6, M. Gauthier3, L. Giuffrida7,10, G. D. Glenn3,8, S. Glenzer3, R. Gray9, J. Green2, G. Hicks6, C. Hyland1, V. Istokskaia7,10, M. King9, D. Margarone1,7, O. McCusker1, P. McKenna9, Z. Najmudin6, C. Parisuana3,11, P. Parsons1, B. Loughran1, C. Spindloe2, D. R. Symes2, A. G. R. Thomas5, F. S. Treffert3,12, N. Xu6, *Influence of novel liquid sheet target on the stability and flux distribution of laser-accelerated proton beams.*. **in preparation.**

DISTRIBUTION A: Distribution approved for public release.

Presentations / invited talks

- Laser and Plasma Accelerator Workshop 2023, Lagos, Portugal: **Alexander G.R. Thomas** and Ryan Sandberg, *Photon Acceleration from Optical to XUV*
- Invited seminar, Instituto Superior Técnico Portugal 2023: **Alexander G.R. Thomas** and Ryan Sandberg, *Photon Acceleration from Optical to XUV*
- IEEE International Conference on Plasma Science (ICOPS) 2022: R. Sandberg; R. Krasny; A. G. R. Thomas *Vlasov Simulation with FARSIGHT*
- APS Division of Plasma Physics Annual Meeting 2021: **Robert Krasny**, Ryan Sandberg and Alexander G.R. Thomas, *The FARRSIGHT Vlasov-Poisson code.*
- APS Division of Plasma Physics Annual Meeting 2021: **Ryan Sandberg** and Alexander G.R. Thomas, *Phase matched plasma wakefield photon acceleration.*
- APS Division of Plasma Physics Annual Meeting 2021: **Stephen E DiIorio**, Benjamin J Winjum, Frank S Tsung, Jennifer Elle, Ricardo A Fonseca, Alexander G Thomas, *A Deterministic Collisional Ionization Module for Particle-in-Cell Codes*
- Invited seminar, Lawrence Berkeley National Laboratories, November 2021: **Alexander G.R. Thomas** and Daniel Seipt, *Modelling Chromatic Emittance Growth in Staged Plasma Wakefield Acceleration to 1 TeV using Nonlinear Transfer Matrices*
- Invited seminar, Lawrence Berkeley National Laboratories, October 2021: **Ryan Sandberg** and Alexander G.R. Thomas, *The FARRSIGHT Vlasov-Poisson code and Phase matched plasma wakefield photon acceleration.*
- 2021 SIAM Conference on Computational Science and Engineering: Ryan Sandberg, Robert Krasny and Alexander G.R. Thomas, *FARRSIGHT:*

DISTRIBUTION A: Distribution approved for public release.

A Forward Adaptively Refined and Regularized Semi-Lagrangian Integral GPU- and Hierarchical Treecode-Accelerated Method for the Vlasov-Poisson System

- APS Division of Plasma Physics Annual Meeting 2020: Stephen DiIorio, Benjamin Winjum, Joshua May, Frank Tsung, Jennifer Elle, Alexander G.R. Thomas, *Towards a Collisional Ionization Model for Particle-in-Cell Codes*

- Michigan Institute for Plasma Science and Engineering Graduate Student Symposium 2020: Ryan Sandberg, Robert Krasny and Alexander G.R. Thomas, *FARRSIGHT: A Forward Adaptively Refined and Regularized Semi-Lagrangian Integral Green's function Hierarchical Treecode accelerated method for the Vlasov-Poisson system*

- APS Division of Plasma Physics Annual Meeting 2019:Diiorio, Stephen; Winjum, Benjamin; May, Joshua; Phillips, Ryan; Elle, Jennifer; Thomas, Alexander, *Modeling Collisional Ionization Using a Modified Binary-Encounter-Bethe Model in the Particle-in-Cell Code OSIRIS*

- APS Division of Plasma Physics Annual Meeting 2019: Sandberg, Ryan; Thomas, Alexander; Krasny, Robert *Numerical finite particle effects and wave breaking limits*

Theses

- Ryan Sandberg, *Vlasov simulation with FARSIGHT and unlimited photon acceleration*, defense completed January 2022.

Software

- FARSIGHT Vlasov code `https://github.com/RTSandberg/FARSIGHT`

- Collisional ionization package for OSIRIS 4.0.

DISTRIBUTION A: Distribution approved for public release.

Impacts

Development of the principal discipline(s) of the project

The computational tools developed on this project will expand plasma simulation capabilities in fidelity and improve modeling of ionizable regimes for relativistic plasma. These have already enabled us to study the physics of laser heated gas density plasma at intermediate intensities and proton beam propagation through water vapor, both relevant to experimental data that was obtained by collaborators on high power laser systems at UM and in the UK.

Describe the impact in this reporting period on the development of human resources

The project supported the dissertation work of two PhD students, Ryan Sandberg and Stephen DiIorio. Sandberg developed the semi-Lagrangian tree code FARSIGHT and studied a new scheme for XUV generation using plasma wakefields. DiIorio has implemented a new collisional ionization model into the OSIRIS 4.0 framework in addition to improvements to the field ionization model. Ryan Sandberg successfully defended his thesis in January 2022. DiIorio is aiming to finish up summer 2023. The project will have supported the development of two PhD students trained in kinetic plasma simulation. DiIorio spent the summer at the AFRL in the High Power Electromagnetics Division, Kirtland AFB under the Airforce summer scholars program, working with Jennifer Elle on the modelling working with computational researchers there in summer 2019.

Describe the impact on teaching and educational experiences

Over the course of the project, the PI developed course material on plasma accelerators that was taught as week-long courses at the USPAS schools on particle accelerators in 2020 and 2022, the latter being remote only. The material was developed into a semester long graduate course on plasma accelerators and photon sources that was taught in 2021 and will be taught again in 2023.

DISTRIBUTION A: Distribution approved for public release.

Describe the impact in this reporting period on physical, institutional, and information resources that form infrastructure.

The project outcomes include new simulation resources, including improved particle-in-cell modeling of partially ionized plasma and a new semi-Lagrangian tree code framework for plasma simulation, FARSIGHT.

Changes

Changes in approach

An additional new study by student Ryan Sandberg was added to the scope, using the plasma wakefield to generate intense XUV radiation. This was successful and led to a paper published in Phys. Rev. Lett.

Problems or delays / Expenditure Impacts

A no-cost extension was requested to mitigate COVID-19 related delays involving the students on the project, and was granted by AFOSR. Owing to restrictions in response to the ongoing COVID-19 pandemic, although the students on the project were able to use the research expenses and stipend funding to carry out computational development and simulations, money allocated for travel was not spent owing to travel restrictions because of COVID-19. Funding allocated for a post-doc was not spent owing to existing post-docs leaving and not being able to be replaced because of hiring restrictions in response to COVID-19. This resulted in several publications that are planned from the research being behind where they would be. Funds will be used to support PI and Co-I, the students, and research related expenses directly related to this project to finish these publications.

DISTRIBUTION A: Distribution approved for public release.

Technical Updates

Project 1: Implementation of new Collisional Ionization module in OSIRIS 4.0

Development and testing of the new Collisional Ionization module is shown below.

1 Introduction

Particle-in-cell (PIC) codes [2, 3, 4] are an essential tool in understanding a wide swath of plasma scenarios. Built upon solving fundamental equations, PIC codes can model plasmas and their evolution with speed and accuracy and oftentimes provide unbeknownst insight into the underlying physical mechanisms of the systems they simulate. Because of this they are used to both nicely complement a plethora of different experiments (ranging from fusion research [5, 6, 7, 8, 9] to wakefield accelerators [10, 11, 12] to even x-ray generation [13, 14, 15]) and explore dynamics in otherwise unobtainable scenarios (think astrophysical plasmas [16, 17, 18, 19, 20, 21]) so that we can better understand our universe.

The success of PIC codes comes from an intricate flow of information between particles and a series of field quantities as they influence one another. Simulation particles (or macro-particles) act as a collective of multiple real-world particles and have positions and velocities defined on a continuum. These particles deposit their information, like how much charge they represent and velocity, onto grids at discrete physical locations. We use these interpolated values to calculate currents; we use these currents to solve for the electromagnetic fields present (again, at discrete mesh locations), and once this data has been accumulated, we interpolate these cumulative grid quantities back onto particles to calculate the forces that move them around. This set of actions is the predominance of a singular simulation timestep. During a simulation, time is advanced at discrete intervals, and this cycle continues for as long as the user desires to simulate their plasma.

In practice, the separation of these physical features into particle and grid quantities are PIC's defining feature and a necessary design choice owing to its success. This back and forth interpolation scheme helps smooth out noise generated from using a small number of simulation particles or a coarse grid and is efficient at capturing the cumulative dynamics within a system.

DISTRIBUTION A: Distribution approved for public release.

Fundamentally, and without much modification, a "standard"[1] PIC code is oftentimes capable of quickly capturing enough physics to properly and accurately grasp the dynamics of interest in a non-collisional plasma. Yet, with the advent of more powerful lasers and the ability to probe deeper into our universe, the possibility of exploring exotic plasmas with increasingly complex physical considerations is only the beginning of a necessity to include more robust models in the tools we use to simulate them.

Collisions (which in the context of plasmas is colloquially understood to mean Coulombic collisions) and additionally collisional ionization, are becoming an evermore important physical phenomena to understand in complex plasma systems. And as of late, the need to relativistically describe these interactions is required as higher energy electrons are able to be accelerated and interact in each of these scenarios.

Collisions and collisional ionization, much like PIC, are equally applicable and interesting in wide swath of scenarios including astrophysical settings [22, 23, 24], fusion [25, 9, 7], ion acceleration [26, 27], ionization-induced defocusing [28], energy transport and stopping power [29, 8], instability modeling [30], solid-density plasmas [31, 32], and cathode, silicon detectors, and rf plasma modeling [33, 34, 35, 36, 37, 38, 39].

In his review, Birdsall [40] considers Burger [33] the first to develop an electron-neutral elastic collision model for use in simulating low-pressure diodes. While designed to simulate only hardsphere collisions, Burger's implementation developed a lot of the groundwork used in later collision implementations for PIC codes. For each particle, they calculated the probability that that particle would experience a collision. They divided the simulation domain into different regions. They iterated over the particles contained within those regions summed their probabilities. Only those particles whose collision probability increased the sum over a value of one experienced a collision. One was subtracted from the running total and the process continued.

Similarly, Shanny *et al.* [41] created a model for colliding electrons with ions and keyed their implementation the Lorentz Gas Model. No energy was transferred between electrons and the stationary ions. Instead, they considered only the cumulative small-angle collisions with ions and calculated a scattering angle for each electron. The electron velocity directions (not

[1]Standard, in this sense, means a code that follows the base PIC loop of solving Maxwell's equations and the Newton-Lorentz equations of motion without other physical considerations.

DISTRIBUTION A: Distribution approved for public release.

magnitudes) were updated each timestep. Randomness was included during the calculation of the angular deflection to match expected distributions and orientations.

As mentioned previously, the interpolation of particle and grid values back-and-forth used in PIC acts to smooth out any short-range particle-particle interactions (short-range being impact parameters that are less than a cell diameter, Δx [40]). As such, once academics started adapting these algorithms for use in PIC, they explicitly added particle-particle interactions to recover any velocity scattering effects.

Hockney and Eastwood [2] developed a hybrid called P^3M [42] that handled long- and short-range particle interactions separately using the usual PIC method and a binary scattering approach, respectively. Takizuka and Abe [43], with further developments from Procassini *et al.* [44], worked on an explicitly momentum conserving binary scattering between electron-electron, electron-ion, and ion-ion particle pairs. Hockney and Eastwood [2] present yet another approach, albeit designed specifically to simulate electrons flowing through a semiconductor lattice. They updated the local electric fields using one (larger) timestep and conducted several Monte Carlo particle collisions (using a modified scattering rate to include null collisions in order to keep it constant [45, 46]) within each if those field-adjusting timesteps. Finally, Birdsall [40] and several collaborators [35, 47, 48, 36, 37] present their own collision scheme. Their algorithm included the ability for scattering, excitation, and ionization interactions between electron collisions with neutrals; scattering and charge exchange events between ion collisions with neutrals; and electron and ion recombination. Based off the particle characteristics (like velocity), they calculate the probability of a collision event occurring and use a Monte Carlo process to determine if one does or does not happen. Additionally, they calculate the cross sections for each of the aforementioned potential interactions they consider and use another random process to decide which of those interactions occur during the collision event.

In a similar vein, several groups developed in parallel collision algorithms where macro-particles interacted with a grid quantity in a very PIC-like fashion. Jones *et al.* [49] (and similarly Lemons *et al.* [50]) deposited the scattering from a collision as a grid quantity and included this force during the pushing portion of the PIC loop. Manheimer *et al.* [51] developed a grid-based Langevin formulation, which is useful in modeling the plasmas with a multi-peaked or very anisotropic electron distribution. Finally, Cohen *et al.* [52] adapted the work of Takizuka and Abe [43] such that particles in-

DISTRIBUTION A: Distribution approved for public release.

teracted with a unique field particle randomly selected from the local cell distribution. These approaches, however, have not been so readily adopted by more widely available PIC codes.

While many early implementations were designed to suit specific applications, pieces from all of these previous methods started to merge together to form what is commonly implemented in PIC today. If not explicit before, macro-particles never interact directly with one another within a PIC cycle but only do so indirectly via calculated currents and the resulting forces they produce. However, the thus far implemented collision modules in PIC codes break this relationship. Instead, macro-particles *do* interact directly as they 'collide' and scatter. During this process, we consider inter- and intra-collisions. In either case, work is done to sort particles and identify which macro-particles are in proximity to others (usually, those identified as being within the same cell are considered local to one another). Once the local particles are identified, in the case of inter-collisions between species A and species B, the particles of one species are randomly sorted before each species' particles are paired together. These pairs then collide, exchange momentum, and scatter. (This works nicely assuming the two species have the same number of macro-particles. If, however, species A has more macro-particles than species B, then the extra particles from A are assigned to random particles in B, effectively meaning that particles in species B may collide multiple times in a single timestep.) In the case of intra-collisions, collisions between macro-particles of the same species, half of the local particles are randomly shuffled before the two halves are paired together to collide. The idea behind this is that after many timesteps, an individual particle will sample the overall distribution of target particles. This pairing process also helps avoid calculating the collisions between each macro-particle and every other macro-particle, which, even if you only consider local particles, scales horribly $\left(\mathcal{O}\left(n^2\right)\right)$ and is computationally inefficient.

While this methodology was developed to simulate coulomb collisions, it is just as applicable when simulating collisional ionization events. As the coulomb collision is calculated for inter-collision particle pairs (electron and different ion species), an ionization event could also be considered for that same pair. To do this, a collisional cross-section is calculated for the electron-ion pair, and a probability based off this cross-section is determined. A random number is generated, and whether or not that number falls within the probability of the event occurring determines if ionization is triggered.

DISTRIBUTION A: Distribution approved for public release.

If an ionization event does occur, then a new macro-particle representing a secondary electron is created and the energy lost by the incident electron and energy transferred to the newly created electron is calculated and distributed. So, in addition to particles scattering and transferring energy from coulomb collisions, ionization events occur using this Monte Carlo process. Pérez *et al.* proposed such an approach [53] (their model is a refinement of that proposed by Nanbu and Yonemura [54] and Sentoku and Kemp [55]) and several PIC codes [56, 57] implemented this.

The Pérez method has become the de facto approach for several PIC implementations, but is not without some shortcomings. Higginson *et al.* [58] pointed out inaccuracies in their statistical methodology when there is a large discrepancy in the weights and energies of the macro-particle pair colliding. The PIC code Smilei [57] uses this corrected version. Morris *et al.* [59] further develop Higginson's corrections by improving the particle pairing process when there is a large discrepancy in the number of macro-particles of the two colliding species. This helps to more accurately estimate the number of ions that are produced during collisional ionization events. This correction is included in the newest versions of the PIC code Epoch [56].

Perhaps a not immediately obvious caveat to any of these approaches is the necessity for a large number of simulation particles to properly model these collisional interactions, especially when the plasma has some thermal distribution (oftentimes $\mathcal{O}\left(10^4\right)$ or more particles per cell [56]). This requirement is to achieve an adequate statistical sampling of the distribution due to the Monte Carlo aspect of this approach. This problem is only exacerbated as more particles are injected into the system as more ionization events occur.

There has been some effort within this infrastructure to try and minimize the amount of work necessary to calculate these interactions. Some PIC codes allow for these collisional interactions to occur every n timesteps instead of every timestep, as is typically done. Additionally, instead of only pairing particles together that are located within the same cell, some codes allow the pairing of particles within a collection of surrounding cells. Because particles have to be randomly sorted during the pairing process, this larger sampling area helps to alleviate the magnitude of the sorting required [57].

This algorithm is fully capable of modeling collisional ionization within a PIC code but not without its shortcomings. Because this approach relies on macro-particles interacting directly with other macro-particles and ionization events occur via a random process, it introduces unwanted noise into the

DISTRIBUTION A: Distribution approved for public release.

simulation and makes the accuracy of this process very dependent on the number of particles used. Here we seek a deterministic algorithm, which will generate a smoother ionization process for low particle statistics.

In this paper, we present a new algorithm for simulating collisional ionization within a PIC framework. Specifically, we steer away from treating these ionization events as particle-particle interactions that occur via a random process but instead go back to transferring information between particles and grids in a very PIC-like fashion. Our method treats and calculates ionization events deterministically as each particles contribution to the ionization rate is calculated explicitly and deposited onto a grid. This grid of ionization rates is then used to advance ion densities, which allows us to track how much new charge is generated each timestep, so we can create newly ionized electrons accordingly. Additionally, the ionization rate grid, with little modification, keeps track of how much energy is lost per cell due to ionization physics. We interpolate this information back onto the particles; this allows for a continuous decrease in the energy of macro-particles as they participate in ionization events and allows for the easy calculation of the new momentum of ionized electrons. Collectively, these particle-to-grid and grid-to-particle information transfers act as a smoothing process, reducing noise considerably compared to other current algorithms. This algorithm is implemented within the PIC code Osiris [60], and we present several simulations highlighting the new physics that is captured when considering collisional ionization in different scenarios.

2 Theory

We first define N_k and B_k to be the electron occupation number and electron binding energy, respectively, of subshell k and T and W to be the kinetic energy of the incident and secondary electron, respectively. We further normalize these terms as follows, defining $w = W/B_k$, $t = T/B_k$, $b' = B_k/m_e c^2$, and $t' = T/m_e c^2$. We finally define $\beta_b = 1-(1+b')^{-2}$ and $\beta_t = 1-(1+t')^{-2}$. We use the differential ionization cross-section presented by Kim *et al.* [61] and modified by Guerra *et al.* [1].

$$\frac{d\sigma_k}{dw} = \sigma_k^0 \left[A_3 f_3(w) + f_2(w) + \frac{2A_2}{t-1} - A_1 f_1(w) \right], \tag{1}$$

DISTRIBUTION A: Distribution approved for public release.

where

$$\begin{aligned}
\sigma_k^0 &= 4\pi a_0^2 \alpha^4 N_k / \left(\beta_t^2 + \chi_{n\ell j}\beta_b^2\right) 2b', \\
f_n(w) &= (w+1)^{-n} + (t-w)^{-n}, \\
A_1 &= \frac{1}{t+1}\frac{1+2t'}{\left(1+t'/2\right)^2}, \\
A_2 &= \frac{t-1}{2}\frac{b'^2}{\left(1+t'/2\right)^2}, \qquad \text{and} \\
A_3 &= \ln\left(\frac{\beta_t^2}{1-\beta_t^2}\right) - \beta_t^2 - \ln\left(2b'\right).
\end{aligned}$$

Here, c is the speed of light in vacuum, m_e is the mass of an electron, α is the fine structure constant, and a_0 is the Bohr radius.

Guerra changed the σ_k^0 term where they reduced the coefficient to only depend on the incident electron kinetic energy and a term that depends only on Z, the atomic number.

$$\begin{aligned}
\chi_{n\ell j} &= \left(\frac{C_{n\ell j}}{B_k}\right) 2R, \\
C_{n\ell j} &= 0.3\frac{Z_{\text{eff}_{n\ell j}}^2}{2n^2} + 0.7\frac{Z_{\text{eff}_{n'\ell' j'}}^2}{2n'^2}, \\
C_{n\ell j_{\text{last}}} &= \frac{Z_{\text{eff}_{n\ell j}}^2}{2n^2}.
\end{aligned}$$

Here, R is the Rydberg energy and Z_{eff} is the effective nuclear charge or atomic number the $n\ell j$ electron experiences (n is the principle quantum number, ℓ is the orbital angular momentum, and j is the total angular momentum). The primed variables stand for the subshell succeeding the unprimed subshell (working from the innermost electron shell to the outermost). The last equation above exists for use with the outermost electrons in the atom since no electrons lie above the valence shell electrons.

We integrate the differential cross-section and sum over all k subshell cross-sections to get the total ionization cross-section for an atom. Since the primary (fast) electron and the secondary (slow) electron are indistinguishable from one another, practically speaking, it is convention to consider any electron with energy less than half the incident particle's kinetic energy (minus the binding energy required to cause ionization) to be the newly ionized

DISTRIBUTION A: Distribution approved for public release.

electron, resulting in the following upper limit of integration.

$$\begin{aligned}\sigma_k &= \int_0^{(t-1)/2} \frac{d\sigma_k}{dw}\, dw \\ &= \sigma_k^0 \left[\frac{A_3}{2}\left(1 - \frac{1}{t^2}\right) + 1 - \frac{1}{t} + A_2 - A_1 \ln t \right] \qquad (2)\end{aligned}$$

$$\bar{\sigma} = \sum_k \sigma_k. \qquad (3)$$

The first moment of the differential ionization cross-section provides the statistical mean of w, which can be used to calculate the average energy lost by the incident electron and gained by the secondary. For completeness, we quote the expressions given in Ref. [53];

$$\begin{aligned}\langle w \rangle_k &= \frac{1}{\sigma_k} \int_0^{(t-1)/2} w \frac{d\sigma_k}{dw}\, dw \\ &= \frac{\sigma_k^0}{\sigma_k} \left\{ \frac{A_3}{2} \frac{(t-1)^2}{t(t+1)} + 2\ln\left(\frac{t+1}{2}\right) - \ln t + A_2 \frac{t-1}{4} \right. \\ &\qquad \left. - A_1 \left[t \ln t - (t+1) \ln\left(\frac{t+1}{2}\right) \right] \right\}, \qquad (4)\end{aligned}$$

which is used to determine the average energy transferred to the secondary electron (note a factor of t present in Eqn. 4 that is missing in the original reference erroneously)

$$\overline{w} = \sum_k \langle w \rangle_k B_k \sigma_k / \bar{\sigma}, \qquad (5)$$

and the average energy lost by the incident electron

$$\bar{\varepsilon} = \sum_k \left[\langle w \rangle_k + 1 \right] B_k \sigma_k / \bar{\sigma}\ . \qquad (6)$$

It is important to note that because σ_k^0 is a function of shell occupancy number, these values need to be calculated for each different charge state of the ion.

DISTRIBUTION A: Distribution approved for public release.

While simpler than its predecessor, equation 3 is still reliant on data specific to the species at hand. In particular, to properly calculate the cross-section for an atom, one must know the electron occupancy numbers of each of its subshells (and the occupancy of each of its ionization states, which can be nontrivial), the binding energies of each of its electrons, and the effective nuclear charge–effective because there are shielding effects from the other electrons present–each electron experiences. Both experimental and computational work has gone into creating reference tables for different atomic binding energies [62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74], screening constants [75, 76, 77, 78], and shell occupation number [66, 69, 70, 74]. With the variety of databases available, work has gone into at least determining which of the binding energy values produce the most accurate replication of experimental cross-section data [79]. From this reference we chose to use the binding energies from Carlson [64], the most recent calculations of screening constants from Guerra *et al.* [78], and electron shell occupation number from the NIST Atomic Spectra Database [74]. These choices used with equation 3 have led to good agreement with other experimental and theoretical cross-section results [80, 81, 82, 83, 84].

3 Implementation

Several codes have developed collisional ionization packages for particle-in-cell codes, for example [53, 56, 57]. These codes use a Monte-Carlo approach to determine whether ionization occurs. Local particles are paired and a random number is generated for every potential collision event, and whether this number is above or below a calculated probability determines whether or not the event occurs. While these codes are able to simulate collisional ionization, the fact that events are based on a random process introduces unwanted noise into the simulation. In fact, trying to properly sample a thermal distribution of particles takes up to several thousands to several tens of thousands of particles-per-cell [56] to achieve results that start to overcome the noise generated by this random process. Here we seek a deterministic algorithm solving the rate equations, which will generate a smoother ionization process for low particle statistics.

Our algorithm does so by explicitly calculating the rate at which collisional ionization should occur throughout the simulation and properly advancing the relevant ion densities while also injecting new macro-particles

DISTRIBUTION A: Distribution approved for public release.

when necessary. More formally, our goal is to determine how the ion number densities, n_i, change over time and to use the amount of charge generated through ionization to inject new macro-particles into the simulation. Doing this requires we find some means of evaluating a set of coupled differential equations that take on the form

$$\frac{dn_i}{dt} = n_{i-1}\mathsf{R}_{i-1} - n_i\mathsf{R}_i, \tag{7}$$

where R_i is the ionization rate per ion for ions in ion state i. In general, for a given ion state of a species, it can either grow in number density as less ionized ions become more ionized or it can diminish in number density as it itself becomes more ionized. Clearly, the outlying cases to this generalization are the neutral atom, which can only decrease in number density (one of our assumptions described below) and the fully ionized ion, which can only increase in number density.

We solve this problem using the unique infrastructure behind ionizable species within OSIRIS. Ionizable species (referred to as 'neutrals' in the OSIRIS framework) are objects that don't contain any macro-particles themselves but contain information about the local number density of every ionization state of that species in every grid cell. A 'neutral' species is nothing more than a collection of grids that describe the density profile throughout the simulation space for each ionization state of that element. The internal structure of this object is important to clarify for later discussions. In the input deck, the user specifies the maximum ionization state a neutral can reach, e.g. we could specify for Helium to be ionized to either a +1 state and then no further or to a maximum of +2; or we could specify for Aluminium to be ionized to a maximum of any charge state between +1 and +13. This maximum considered ionization state will be referred to as $z_{\max}$. We create $z_{\max} + 1$ grids (+1 to account for a charge state of 0), and each grid cell contains a value between 0 and 1. The value of a cell in essence captures the percentage of the number density of the neutral object that is in that particular ionization state. During the set up of a simulation, every cell in the $\overline{Z}_0$ grid (the grid for charge state 0 or neutrals) is initialized with a value of 1, and every other cell in the remaining grids are assigned 0. As the program proceeds, these values increase and decrease and flow from one grid to the next as we ionize particles.

There is an obvious constraint with an infrastructure designed this way. The first is that the sum of the values in a cell with the same index, j, in

DISTRIBUTION A: Distribution approved for public release.

each of these grids must be 1. In other words

$$\sum_{i=0}^{z_{\max}} \overline{Z}_{i,j} = 1. \tag{8}$$

Second, as we have presented this object so far, we would be unable to fully reconstruct a complete profile of the neutral. We utilize two more grids to rectify this. The first is a grid describing the profile of the neutral object, $\overline{Z}_{\text{PROF}}$. The user specifies a profile they want their species to take on, and that profile is discretized onto this grid. Cell values in this case are ≥ 0. Subsequently, to get the number density for any given charge state of the ion at any timestep, we can calculate the Hadamard product and output

$$n_i = \overline{Z}_i \circ \overline{Z}_{\text{PROF}}. \tag{9}$$

The final grid we create keeps track of how ionized the neutral is at each grid point and is limited to values between 0 and $z_{\max}$. The amount of charge in each cell is calculated every timestep as follows:

$$\overline{Z}^* = \sum_{i=1}^{z_{\max}} i\overline{Z}_i. \tag{10}$$

One inherit limitation with this design choice is that neutral atoms and its associated ions are not capable of moving. An object does exist currently within OSIRIS that allows for moving ions, but it is limited to injecting ions that have been ionized only once, and once they have been ionized they can no longer partake in any ionization process. We tackle this limitation later but, for now, assume that ions are immobile.

Associated with a neutral is a separate species (i.e. representing electrons) that stores information about the secondary particles that get generated during an ionization event. The fact that the ionization state information is stored on a grid allows for the deterministic advancement of each ionization state's number density using the collisional ionization rates.

Our algorithm can be divided into three distinct processes that take place over a simulation timestep:

1. calculation of the collisional ionization rate and contributions to momentum lost/transferred for each specified macro-particle species on a spatial grid,

DISTRIBUTION A: Distribution approved for public release.

2. deterministic advancement of the local ion densities using the calculated rates,

3. injection of new macro-particles if enough charge has been generated with the proper handling of any momentum transfer.

There are a few assumptions we make within our algorithm. First, we calculate the cross-sections for each ion assuming that the ion is in its ground state. We do this in part because the OSIRIS infrastructure does not readily record the information necessary to keep track of all the possible excitation states without significant modifications to the core of the code, well beyond the purview of this project. Additionally the amount of overhead necessary to implement something of this complexity does not currently, we feel, outweigh any potential loss of physical accuracy when the lifetimes of these excited states (hundreds of picoseconds) is well beyond the timescales PIC codes are usually used to study (a few picoseconds if not usually femtosecond timescales). We also calculate these cross-sections assuming the ions are ionized in order from their outermost electron to their innermost. In other words, we never ionize an electron in a 1s shell over one in the current valence shell. We justify this assumption, in part, since the electrons in the outermost shells have a larger cross-section (sometimes orders of magnitude larger) than the inner electrons. Thus, if an interaction were to occur, it would most likely occur by an interaction between the incident electron and one of the outer electrons.

Second, we ignore any three body recombination, which will occur on timescales that are typically much longer than the PIC simulation timescale or electron excitation and deexcitation. Third, rates are calculated ignoring ionization due to ion impact. Finally, when new particles are injected due to collisional ionization, they are all injected with the same calculated average momentum. Additionally, momentum is removed from current simulation particles that participate in ionization regardless if a new particle is generated that timestep or not. This is described in more detail below.

3.1 Initialization of Data Table

During the initialization of the PIC simulation, we generate a data table containing the ionization cross-section (Eqn. 3), energy transferred to the secondary electron (Eqn. 5), and energy lost by the incident electron (Eqn. 6)

DISTRIBUTION A: Distribution approved for public release.

over a sample of incident electron energies evenly spaced on a logarithmic scale. The sampled energies are calculated as follows:

$$1 \tag{11}$$

where n_p is the number of points to sample the data table with, and $e_{\min}/e_{\max}$ are the min/max energies to sample the data table over. This allows us to sample the data over large energy scales (e.g. $\sim$eV to $\sim$TeV) that might be present within any given simulation using only a few hundred sample points. We calculate these functions using energy $(\gamma - 1)$ because we achieve a better sampling at lower energies as opposed to using a variable like γ, where a considerable portion of the values we sample are around 1. This even sampling also allows us to directly calculate the nearest index of a data entry when we later interpolate the data to a continuous macro-particle energy, removing the need to do any searching through the database.

3.2 Collisional Ionization Rate

Three core aspects to the algorithm occur at this point in a timestep. We calculate the rate of ionization occurring, how much momentum the incident particles lose, and how much momentum we will transfer to newly generated particles. All three of these calculations require information from the current simulation particles, so, motivated by overall code efficiency, we handle them all concurrently as we iterate over all simulation particles once. We break down how each of the calculations are done.

The rate for collisional ionization for a particular ion, R_i, is

$$\mathsf{R}_i \equiv n_e \langle \bar{\sigma}_i v \rangle, \qquad \text{where} \tag{12}$$
$$n_e \langle \bar{\sigma}_i v \rangle = \sum_k \iiint_\Omega \sigma_{i,k}(v) v f_e(\mathbf{p}) d^3\mathbf{p}.$$

Here, n_e is the electron number density, v is the velocity of the particle whose ionization rate we are depositing, $\bar{\sigma}_i$ is the cumulative cross section for the ion species, and $f_e(\mathbf{p})$ is the momentum distribution function of the ionizing particles.

Instead of simulating collisional ionization through particle-particle interactions, like most other PIC codes do leading to the need of some Monte Carlo sampling process, we instead exactly calculate the rate at which collisional ionization should occur given information about the macro-particles at any given timestep.

DISTRIBUTION A: Distribution approved for public release.

For each neutral species, a grid is first created for each ionization state with the same mesh structure as used for the EMF field. We iterate over every particle in each collisionally ionizing species that the user specifies and deposit their contribution to the total ionization rate onto the grid. In practice, to evaluate equation 12 and deposit the ionization rate, we iterate over every particle and deposit their particle weight, velocity, and cross-section onto the ionization rate grid for each ionization level and distribute this quantity according to the particle's shape function. The total ionization rate, R, for ion species i is therefore calculated at time step m through the summation over P particles

$$\mathsf{R}_{i,j}^{m} = \sum_{p=1}^{P} S\big(\mathbf{X}_j - \mathbf{x}_p\big)\big|q_p\big|\overline{\sigma_i}\big(\gamma_p\big)\big|\mathbf{v}_p\big|, \tag{13}$$

where $S\big(\mathbf{X}_j - \mathbf{x}_p\big)$ is the particle shape factor at grid cell j with position $\mathbf{X}_j$ and particle position $\mathbf{x}_p$, $\big|q_p\big|$ is the particle's weight (or equivalently the absolute value of charge in OSIRIS), $\overline{\sigma_i}\big(\gamma_p\big)$ is the cross section of the ion species for an electron with energy γ_p, and $\big|\mathbf{v}_p\big|$ is the speed of the particle.

At the same time as this rate is being calculated, we additionally remove energy from the incident particles contributing to the ionization rate and calculate the amount of energy being transferred to newly created particles. We gain inspiration from SMILEI [57] in calculating the following changes in momentum. Since one limitation of our approach is that ionizable ions are stationary, our lab frame is inherently the ion rest frame, making some of our expressions simpler than their counterparts presented by SMILEI.

As we iterate over every particle to calculate the rate of ionization, we actively keep track of the change in all components of that particle's momentum, $\mathbf{p}$, as it causes ionization with each present ion state for a given neutral. A singular particle's, p, change in momentum is calculated as

$$\Delta\mathbf{p}_p = \sum_{i,j} S\big(\mathbf{X}_j - \mathbf{x}_p\big)\overline{\sigma_i}\big(\gamma_p\big)\big|\mathbf{v}_p\big|n_{ij}\Delta t(\alpha_\varepsilon - 1)\mathbf{p}_p, \qquad \text{where} \tag{14}$$

$$\alpha_\varepsilon = \frac{\sqrt{(\gamma_p - \bar{\varepsilon})^2 - 1}}{\sqrt{\gamma_p^2 - 1}}. \tag{15}$$

Then, once we are finished calculating that particle's contribution to the ionization rate, we update its momentum simply by $\mathbf{p}'_p = \mathbf{p}_p + \Delta\mathbf{p}_p$. One

DISTRIBUTION A: Distribution approved for public release.

item of note is that there is no inherent scattering of the particles included in this calculation that would change $\hat{\mathbf{p}}_p$.

Finally, at this point in the algorithm we would like to know how much momentum to transfer to newly ionized electrons due to these ionization processes. However, newly ionized macro-particles do not exist yet within the simulation. So, we generate one final set of grids for each momentum component of similar dimensions and spacing to the ionization grid that tracks how much momentum we would have transferred to a new particle during each ionization event. We first deposit the contribution of momentum transferred, $\mathbf{p}_w$, by each incident particle onto the grid via the following means

$$\mathbf{p}^m_{w\,j} = \sum_{p=1}^{P} \sum_i S\big(\mathbf{X}_j - \mathbf{x}_p\big) \big|q_p\big| \overline{\sigma_i}\big(\gamma_p\big) \big|\mathbf{v}_p\big| n_{ij} \Delta t \alpha_w \mathbf{p}_p, \qquad \text{where} \tag{16}$$

$$\alpha_w = \frac{\sqrt{(\overline{w}+1)^2 - 1}}{\sqrt{\gamma_e^2 - 1}}. \tag{17}$$

Once the contributions of every particle are accounted for, we calculate the average momentum to be transferred to newly injected particles.

$$\langle \mathbf{p}_w \rangle^m_j = \frac{\mathbf{p}^m_{w\,j}}{\sum_{p=1}^{P} \sum_i S\big(\mathbf{X}_j - \mathbf{x}_p\big) \big|q_p\big| \overline{\sigma_i}\big(\gamma_p\big) \big|\mathbf{v}_p\big| n_{ij} \Delta t} \tag{18}$$

We store this information and use it later during the particle injection phase of the algorithm, which we describe below.

Recall that at the initialization step of the program we populate a data table with calculated values for the cross-section, energy transferred, and energy lost for an ion species. In all of the previously mentioned calculations, we make use of this data table by interpolating and extrapolating values from this table rather than explicitly evaluating equations 3, 5, and 6 for every particle at every timestep. The kinetic energy, T, of each particle is calculated and then used to find the correct index in the data table to interpolate the data between without the need to do any searching. Because we sampled our functions at regular energy intervals, we calculate this index, i, using

$$\begin{aligned} x &= \log_{10}\big(T/e_{\min}\big) \frac{n_p - 1}{\log_{10}\big(e_{\max}/e_{\min}\big)} + 1 \\ i &= \lfloor x \rfloor \end{aligned}$$

DISTRIBUTION A: Distribution approved for public release.

(the +1 is omitted if using a programming language that has zero indexed arrays). If the energy of the particle is below or above $e_{\min}$ or $e_{\max}$, we set $i = 1$ or $i = n_p - 1$ (0 and $n_p - 2$ for zero indexed languages), respectively, which gives us two data points to extrapolate values outside the range of our data table. Using our cross-section table as an example, we then perform a linear interpolation to gather the desired value, $\bar{\sigma}'$, via

$$\bar{\sigma}' = \big(\bar{\sigma}[i+1] - \bar{\sigma}[i]\big)a + \bar{\sigma}[i], \qquad \text{where}$$
$$a = x - i.$$

3.3 Ion Density Advancement

Once the ionization rate has been calculated, we can deterministically advance the densities of each ionization state of the neutral object. We aim at this point to evaluate equation 7. Given a rate of ionization, R, the state densities, n_i, at every grid point can be advanced from the m^{th} timestep to the next using a second order Runge-Kutta method. However, given the internal structure of the neutral object, discussed in length above, we do not require the calculation of n_i to proceed, and we can capture the same information

DISTRIBUTION A: Distribution approved for public release.

using the $\overline{Z}_i$ grids.

$$\begin{aligned}
\frac{dn_i}{dt} &= n_{i-1}\mathsf{R}_{i-1} - n_i\mathsf{R}_i \\
\\
\frac{dn_0}{dt} &= -n_0\mathsf{R}_0 \\
\frac{dn_1}{dt} &= n_0\mathsf{R}_0 - n_1\mathsf{R}_1 \\
&\vdots \\
\frac{dn_{z_{\max}}}{dt} &= n_{z_{\max}-1}\mathsf{R}_{z_{\max}-1}
\end{aligned}$$

$$\left.\begin{aligned}
\implies \overline{Z}_0^{m+1} &= \overline{Z}_0^m + \left[-\overline{Z}_0^m\mathsf{R}_0\Delta t\left(1+\frac{1}{2}\mathsf{R}_0\Delta t\right)\right] \\
\overline{Z}_1^{m+1} &= \overline{Z}_1^m + \left[\overline{Z}_0^m\mathsf{R}_0\Delta t\left(1+\frac{1}{2}\mathsf{R}_0\Delta t\right) - \overline{Z}_1^m\mathsf{R}_1\Delta t\left(1+\frac{1}{2}\mathsf{R}_1\Delta t\right)\right] \\
&\vdots \\
\overline{Z}_{z_{\max}}^{m+1} &= \overline{Z}_{z_{\max}}^m + \left[\overline{Z}_{z_{\max}-1}^m\mathsf{R}_{z_{\max}-1}\Delta t\left(1+\frac{1}{2}\mathsf{R}_{z_{\max}-1}\Delta t\right)\right]
\end{aligned}\right\} \tag{19}$$

Implicit in this formulation is that this is calculated per grid cell.

3.4 Macro-particle Injection

After advancing the densities, particle injection is determined and carried out on a cell by cell basis. For each cell, we compare how much the charge density has changed from the last timestep to the current, and if the cumulative charge generated from ionizing each of the states is great enough, we create and inject new macro particles into the current grid cell. The number of particles to inject into a cell is determined via

$$n_{\text{add}} = \left\lfloor \frac{\overline{Z}_m^*}{z_{\max}} n_{\text{ppc}} + 0.5 \right\rfloor - \left\lfloor \frac{\overline{Z}_{m-1}^*}{z_{\max}} n_{\text{ppc}} + 0.5 \right\rfloor \tag{20}$$

DISTRIBUTION A: Distribution approved for public release.

For every particle injected into a cell, j, we give them a weight of

$$q = -\frac{z_{\max}\overline{Z}_{\text{PROF},j}}{n_{\text{ppc}}}. \tag{21}$$

If field ionization also took place, we proportionally divide how many particles to inject into the simulation due to field ionization and collisional ionization by their corresponding ionization rates. We give particles injected with field ionization an initial momentum of zero. But, specifically for the particles injected via collisional ionization, we have to calculate a momentum transfer that took place since we are directly simulating the action of two particles colliding.

We iterate through every simulation cell, and if enough charge has been generated to warrant the injection of new particles from collisional ionization we make use of the calculated average momentum transferred we stored earlier from Eqn. 18 to assign those particles a momentum. Following the PIC-like paradigm, the momentum transferred is interpolated from the grid we previously calculated onto the to-be-injected particle's position and assigned to that particle. A new particle, p', will be assigned a momentum given by

$$\mathbf{p}_{p'} = \sum_j \langle \mathbf{p}_w \rangle_j^m \, S\big(\mathbf{X}_j - \mathbf{x}_{p'}\big). \tag{22}$$

3.5 Comments on Algorithm

(Mention how this entire paradigm can be generalized and applied to any type of ionization or rate changing physics)

4 Collisional Ionization with Ionizable Moving Ions

Our entire discussion prior relied on neutral values being stored as grid quantities, but this resulted in the ions themselves being immobile. This is acceptable in some situations, e.g. simulations on short enough timescales where ion motion does not matter or simulations with a moving window where ions leave the simulation box relatively quickly anyway. But this becomes a hindrance more than a blessing once we care about ion dynamics in response

DISTRIBUTION A: Distribution approved for public release.

to a plasma's evolution. Our next goal is to discuss a framework that does allow for ionizable moving ions while reusing as many of our previous tools developed as we can.

The first, and probably most significant, change made is how we structure the ionization data. Because we desire for macro-particles that can both move and ionize, we move away from storing how far along the object is ionized from a grid value to making them a particle value. Instead, we created a new variable for each species object, ζ, that tracks how far along each *particle* is ionized between the species' current charge state to its next. If z^*, a continuous value between 0 and $z_{\max}$, is the ionization state of the particle, then we define

$$\zeta = 1 - \left(z^* - \lfloor z^* \rfloor\right). \tag{23}$$

At simulation initialization, we create $z_{\max} + 1$ species that are associated with the neutral object. Each species individually describes a different ionization state for that neutral within the simulation (the neutral state all the way up to the max ionization state desired). During setup, we assign only the species for charge state 0 macro-particles to correctly build the density profile desired. These particles are all assigned a value of $\zeta = 1$. The rest of the species have no macro-particles at this point. In addition to every particle being assigned a ζ value, we also assign every particle a random number between (0,1) (non-inclusive). This random number acts as a promotion-threshold for that particular particle, and once its ζ value reaches that criteria, we promote that particle to the species corresponding to the next ionization state.

Just as before, we still need to calculate the ionization rate throughout the simulation. We still make use of the rate grids, R_i, and deposit the contributions to the ionization rate from each participating electron species as described in section 3.2. The removal of momentum from participating electrons and the generation of the transferred momentum grids are calculated similarly as before. One nicety of this setup is that we can use a similar description for how the system evolves as we did before. As a simulation progresses and ionization occurs, the value of ζ for every particle decreases according to the same rate equations we already developed:

$$\frac{d\zeta_i}{dt} = \cancel{\zeta_{i-1} R_{i-1}} \xrightarrow{} 0 - \zeta_i R_i \tag{24}$$

$$\zeta_i^{m+1} = \zeta_i^m - \zeta_i^m R_i \Delta t 1 + \frac{1}{2} R_i \Delta t. \tag{25}$$

DISTRIBUTION A: Distribution approved for public release.

The ionization rate, still a grid quantity, is interpolated onto each macro-particle as it experiences the cumulative effects from all participating electrons. Once the value of ζ for a given macro-particle reaches its promotion threshold we perform a particle promotion. Upon a promotion, we create a new particle with the same position and momentum as the old one, and we update the particle's charge (synonymous with weight) by a factor of z_{i+1}/z_i. The ζ value of the new ion macro-particle is reset back to a value of 1, and a new random promotion-threshold is generated for that macro-particle. Once the new particle is created, the old particle is deleted from the simulation.

In addition to the creation of the new ion macro-particle, a new electron macro-particle is also generated within the same cell as the ion but with a randomized position within the cell and is assigned a charge of $-q_i/z_i$, where q_i is the charge of the ion macro-particle being promoted. The position of the electron is randomized so that the ion and electron are not created at the same location where their current contributions would cancel each other out, which essentially removes their contribution to any of the plasma dynamics. (This is less of an issue at higher ionization stages where the charges of either particle do not exactly cancel out, but is an issue for the first ionization state.) The momentum of the new electron is also handled just as it was in section 3.4.

A stark difference to the prior implementation is how we balance the inflow and outflow of number density expressed with equation 24. As it is currently implemented, the value of ζ only ever decreases as a macro-particle becomes more "ionized." Even if macro-particles representing a lower ionization state see a decrease in their ζ value, macro-particles in higher ionization states never see an increase, hence why the we set the first term in equation 24 to zero. Instead, the balance of number density between these different species is handled completely through the creation and deletion process as macro-particles get promoted. This design choice proves to be sufficient in matching the numerical solutions to the differential equations describing the changes in number density.

DISTRIBUTION A: Distribution approved for public release.

5 Benchmarking

5.1 Collisional Ionization Rates

Arguably the most crucial portion of the algorithm to get correct is the rate at which ionization occurs since every succeeding aspect of the program depends on its accuracy.

5.2 Momentum Loss/Transfer

5.2.1 Momentum Loss

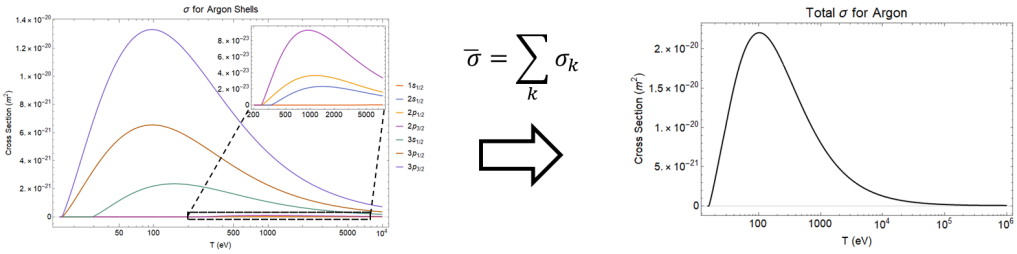


Figure 1: To calculate the cross section for an ion we first calculate the cross section of each occupied subshell using equation 3. Then the total cross section is given by summing all of the individual subshell cross sections.

Algorithm 1: Impact Ionization Rates

for *every particle, i,* **do**
  for *every ionization state, s,* **do**
    R ←[ ] `// grid to store calculated ionization rates`
    KE_i ←kinetic energy of particle i
    $\sigma_{i,s}$ ←`InterpolateCrossSection` (s, KE_i) `// interpolate from look-up table`
    $\mathsf{r}_i \leftarrow w_i \sigma_{i,s} \|\vec{p}_i\|_2 / \gamma_i$ `// calculate rate`
    R ←`DepositOnGrid` (r_i) `// deposit rate on grid based on particle shape function`
  end
end

DISTRIBUTION A: Distribution approved for public release.

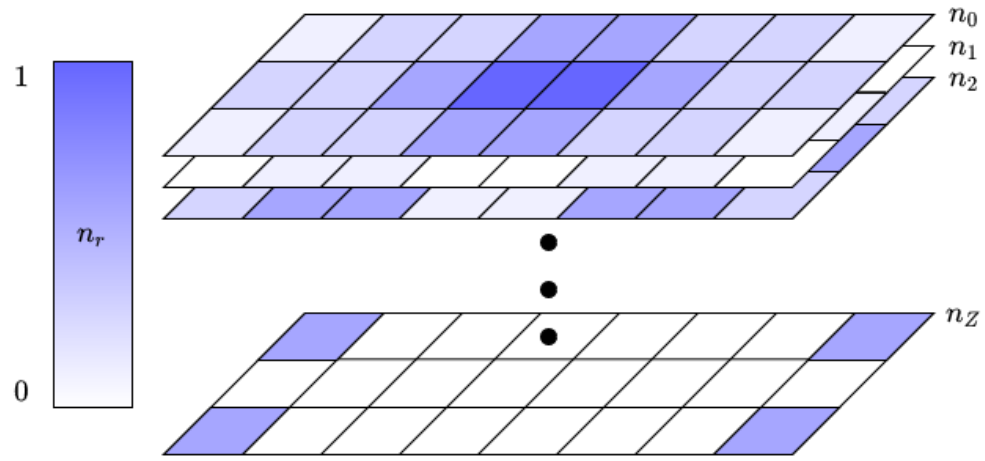


Figure 2: Diagram showing the neutral species' data structure. Ion densities are stored on a grid the same size as the simulation grid for the neutral densities (n_0) up to a full ion (n_Z). For demonstrative purposes, the densities shown are all in relation to some reference density, n_r. It is also worth noting that densities within cells of the same index across multiple ionization states must sum up to the reference density.

The new algorithm is tested against an analytic stopping power calculation of Rohrlich and Carlson and also simulations with the Monte-Carlo particle-in-cell code Smilei in Fig. 5, showing good agreement. Fig. 6 shows the convergence of the algorithm with the number of particles per cell (PPC), showing that few PPC are needed to get a reasonable ionized electron distribution.

References

[1] M. Guerra, F. Parente, P. Indelicato, and J.P. Santos. Modified binary encounter Bethe model for electron-impact ionization. *International Journal of Mass Spectrometry*, 313:1–7, March 2012.

[2] Roger W. Hockney and James W. Eastwood. *Computer Simulation Using Particles*. McGraw-Hill International Book Co, New York, 1981.

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

DISTRIBUTION A: Distribution approved for public release.

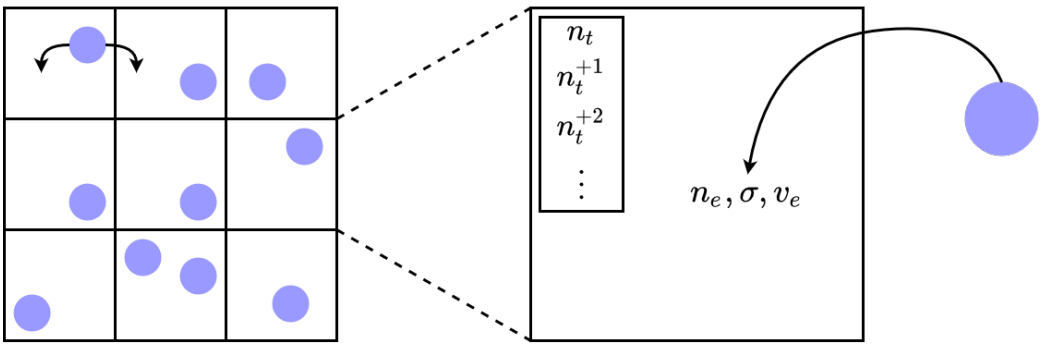


Figure 3: Particle quantities like weight (density or charge) and speed are collected from all of the simulation particles. In addition to this, the collisional ionization cross section for each of the different ion levels is calculated using the incident particle's speed. These three values are used to determine the collisional ionization rate for each ion level in the simulation. These values are interpolated onto different grids for each ion level. To determine the total amount of ionization that should occur throughout the simulation, all of these rates are deposited from all of the particles.

[4] Charles K. Birdsall and A. Bruce Langdon. *Plasma Physics via Computer Simulation*. McGraw-Hill, New York, 1985.

[5] Susanne Pfalzner. *An Introduction to Inertial Confinement Fusion*. CRC Press, March 2006.

[6] C. Bellei, L. Divol, A. J. Kemp, M. H. Key, D. J. Larson, D. J. Strozzi, M. M. Marinak, M. Tabak, and P. K. Patel. Fast ignition: Dependence of the ignition energy on source and target parameters for particle-in-cell-modelled energy and angular distributions of the fast electrons. *Physics of Plasmas*, 20(5):052704, May 2013.

[7] A.P.L. Robinson, D.J. Strozzi, J.R. Davies, L. Gremillet, J.J. Honrubia, T. Johzaki, R.J. Kingham, M. Sherlock, and A.A. Solodov. Theory of fast electron transport for fast ignition. *Nuclear Fusion*, 54(5):054003, May 2014.

[8] P. Norreys, D Batani, S Baton, F N. Beg, R. Kodama, P.M. Nilson, P. Patel, F. Pérez, J.J. Santos, R.H.H. Scott, V.T. Tikhonchuk, M. Wei,

DISTRIBUTION A: Distribution approved for public release.

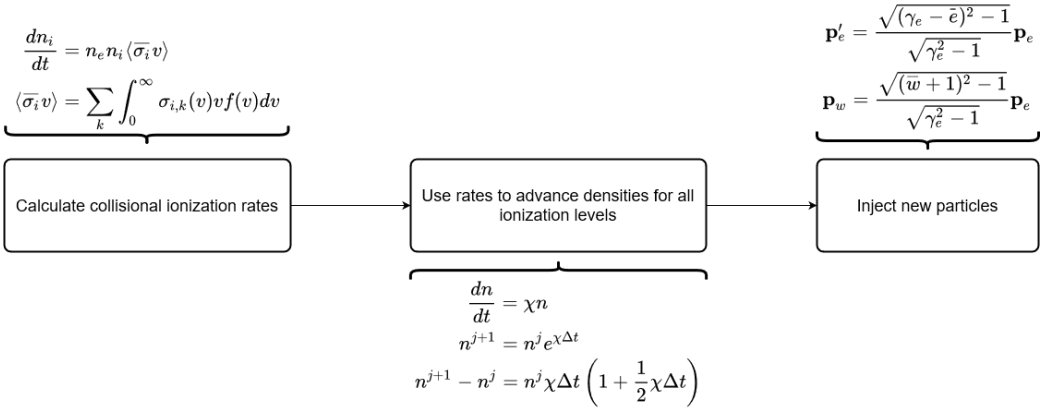


Figure 4: An overview of the collisional ionization algorithm, highlighting its three key components and their related equations.

and J. Zhang. Fast electron energy transport in solid density and compressed plasma. *Nuclear Fusion*, 54(5):054004, May 2014.

[9] A.J. Kemp, F. Fiuza, A. Debayle, T. Johzaki, W.B. Mori, P.K. Patel, Y. Sentoku, and L.O. Silva. Laser–plasma interactions for fast ignition. *Nuclear Fusion*, 54(5):054002, May 2014.

[10] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J.-P. Rousseau, F. Burgy, and V. Malka. A laser–plasma accelerator producing monoenergetic electron beams. *Nature*, 431(7008):541–544, September 2004.

[11] C. G. R. Geddes, Cs Toth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans. High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding. *Nature*, 431(7008):538–541, September 2004.

[12] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick. Monoenergetic beams of relativistic electrons from intense laser–plasma interactions. *Nature*, 431(7008):535–538, September 2004.

DISTRIBUTION A: Distribution approved for public release.

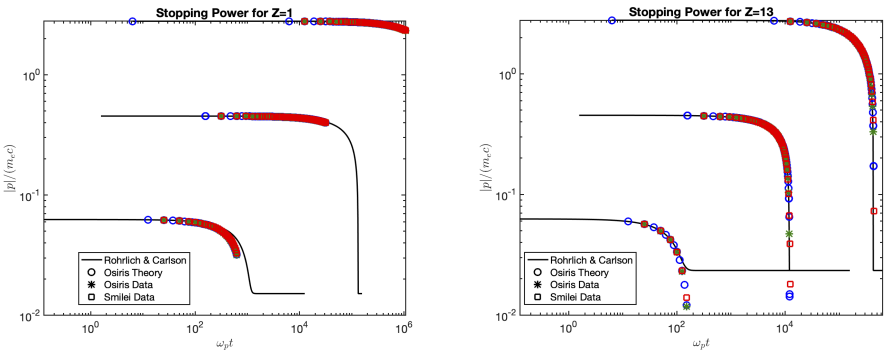


Figure 5: Stopping power from new model compared with Monte-Carlo particle-in-cell code Smilei and analytic stopping power calculation of Rohrlich and Carlson. (Left) Stopping power for $Z = 1$, (right) Stopping power for $Z = 13$.

[13] I. Kostyukov, S. Kiselev, and A. Pukhov. X-ray generation in an ion channel. *Physics of Plasmas*, 10(12):4818–4828, December 2003.

[14] Antoine Rousse, Kim Ta Phuoc, Rahul Shah, Alexander Pukhov, Eric Lefebvre, Victor Malka, Sergey Kiselev, Fréderic Burgy, Jean-Philippe Rousseau, Donald Umstadter, and Daniéle Hulin. Production of a keV X-Ray Beam from Synchrotron Radiation in Relativistic Laser-Plasma Interaction. *Physical Review Letters*, 93(13):135005, September 2004.

[15] F. V. Hartemann, D. J. Gibson, W. J. Brown, A. Rousse, K. Ta Phuoc, V. Mallka, J. Faure, and A. Pukhov. Compton scattering x-ray sources driven by laser wakefield acceleration. *Physical Review Special Topics - Accelerators and Beams*, 10(1):011301, January 2007.

[16] Tsvi Piran. The physics of gamma-ray bursts. *Reviews of Modern Physics*, 76(4):1143–1210, January 2005.

[17] Mikhail V. Medvedev, Massimiliano Fiore, Ricardo A. Fonseca, Luis O. Silva, and Warren B. Mori. Long-Time Evolution of Magnetic Fields in Relativistic Gamma-Ray Burst Shocks. *The Astrophysical Journal*, 618(2):L75–L78, January 2005.

[18] E Waxman. Gamma-ray bursts and collisionless shocks. *Plasma Physics and Controlled Fusion*, 48(12B):B137–B151, December 2006.

DISTRIBUTION A: Distribution approved for public release.

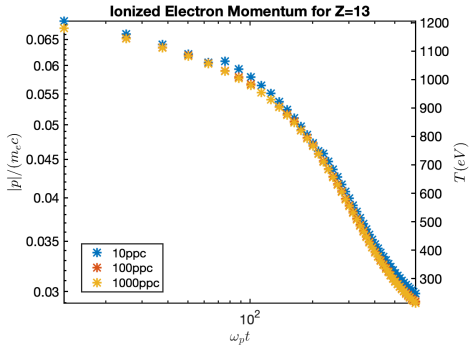


Figure 6: Convergence of the distribution of ionized electrons with number of particles per cell.

[19] F. Fiuza, R. A. Fonseca, J. Tonge, W. B. Mori, and L. O. Silva. Weibel-Instability-Mediated Collisionless Shocks in the Laboratory with Ultraintense Lasers. *Physical Review Letters*, 108(23):235004, June 2012.

[20] L. Sironi, U. Keshet, and M. Lemoine. Relativistic Shocks: Particle Acceleration and Magnetization. *Space Science Reviews*, 191(1-4):519–544, October 2015.

[21] Lorenzo Sironi, Dimitrios Giannios, and Maria Petropoulou. Plasmoids in relativistic reconnection, from birth to adulthood: First they grow, then they go. *Monthly Notices of the Royal Astronomical Society*, 462(1):48–74, October 2016.

[22] Ralph S. Sutherland and M. A. Dopita. Cooling Functions for Low-Density Astrophysical Plasmas. *The Astrophysical Journal Supplement Series*, 88:253, September 1993.

[23] M. Arnould, S. Goriely, and K. Takahashi. The r-process of stellar nucleosynthesis: Astrophysics and nuclear physics achievements and mysteries. *Physics Reports*, 450(4-6):97–213, September 2007.

[24] Alexandre Marcowith, Gilles Ferrand, Mickael Grech, Zakaria Meliani, Illya Plotnikov, and Rolf Walder. Multi-scale simulations of particle acceleration in astrophysical systems. *Living Reviews in Computational Astrophysics*, 6(1):1, December 2020.

DISTRIBUTION A: Distribution approved for public release.

[25] A. A. Harms, K. F. Schoepf, G. H. Miley, and D. R. Kingdon. *Principles of Fusion Energy: An Introduction to Fusion Energy for Students of Science and Engineering*. World Scientific, 2008.

[26] Matthew Allen, Pravesh K. Patel, Andrew Mackinnon, Dwight Price, Scott Wilks, and Edward Morse. Direct Experimental Evidence of Back-Surface Ion Acceleration from Laser-Irradiated Gold Foils. *Physical Review Letters*, 93(26):265004, December 2004.

[27] M. Afshari, S. Morris, L. D. Geulig, Z. M. Chitgar, P. Gibbon, P. G. Thirolf, and J. Schreiber. The role of collisional ionization in heavy ion acceleration by high intensity laser pulses. *Scientific Reports*, 12(1):18260, October 2022.

[28] S.C. Rae. Ionization-induced defocusing of intense laser pulses in high-pressure gases. *Optics Communications*, 97(1-2):25–28, March 1993.

[29] Thomas Peter and Jürgen Meyer-ter-Vehn. Energy loss of heavy ions in dense plasma. II. Nonequilibrium charge states and stopping powers. *Physical Review A*, 43(4):2015–2030, February 1991.

[30] C. Ruyer, S. Bolaños, B. Albertazzi, S. N. Chen, P. Antici, J. Böker, V. Dervieux, L. Lancia, M. Nakatsutsumi, L. Romagnani, R. Shepherd, M. Swantusch, M. Borghesi, O. Willi, H. Pépin, M. Starodubtsev, M. Grech, C. Riconda, L. Gremillet, and J. Fuchs. Growth of concomitant laser-driven collisionless and resistive electron filamentation instabilities over large spatiotemporal scales. *Nature Physics*, 16(9):983–988, September 2020.

[31] D. Wu, X. T. He, W. Yu, and S. Fritzsche. Monte Carlo approach to calculate ionization dynamics of hot solid-density plasmas within particle-in-cell simulations. *Physical Review E*, 95(2):023208, February 2017.

[32] D. Wu, X. T. He, W. Yu, and S. Fritzsche. Particle-in-cell simulations of laser–plasma interactions at solid densities and relativistic intensities: The role of atomic processes. *High Power Laser Science and Engineering*, 6:e50, 2018.

[33] Peter Burger. Elastic Collisions in Simulating One-Dimensional Plasma Diodes on the Computer. *Physics of Fluids*, 10(3):658, 1967.

DISTRIBUTION A: Distribution approved for public release.

[34] J P Boeuf and E Marode. A Monte Carlo analysis of an electron swarm in a nonuniform field: The cathode region of a glow discharge in helium. *Journal of Physics D: Applied Physics*, 15(11):2169–2187, November 1982.

[35] R. W. Boswell and I. J. Morey. Self-consistent simulation of a parallel-plate rf discharge. *Applied Physics Letters*, 52(1):21–23, January 1988.

[36] David Vender. *Numerical Studies of the Low Pressure RF Plasma*. PhD thesis, The Australian National University, Canberra, September 1990.

[37] D. Vender and R.W. Boswell. Numerical modeling of low-pressure RF plasmas. *IEEE Transactions on Plasma Science*, 18(4):725–732, August 1990.

[38] V. Vahedi and M. Surendra. A Monte Carlo collision model for the particle-in-cell method: Applications to argon and oxygen discharges. *Computer Physics Communications*, 87(1-2):179–198, May 1995.

[39] Hans Bichsel. A method to improve tracking and particle identification in TPCs and silicon detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 562(1):154–197, June 2006.

[40] C.K. Birdsall. Particle-in-cell charged-particle simulations, plus Monte Carlo collisions with neutral atoms, PIC-MCC. *IEEE Transactions on Plasma Science*, 19(2):65–85, April 1991.

[41] Ramy Shanny, John M. Dawson, and John M. Greene. One-Dimensional Model of a Lorentz Plasma. *Physics of Fluids*, 10(6):1281, 1967.

[42] J.W. Eastwood, R.W. Hockney, and D.N. Lawrence. P3M3DP—The three-dimensional periodic particle-particle/ particle-mesh program. *Computer Physics Communications*, 19(2):215–261, April 1980.

[43] Tomonor Takizuka and Hirotada Abe. A binary collision model for plasma simulation with a particle code. *Journal of Computational Physics*, 25(3):205–219, November 1977.

[44] R J Procassini, C K Birdsall, E C Morse, and B I Cohen. A relativistic Monte Carlo binary collision model for use in plasma particle simulation

DISTRIBUTION A: Distribution approved for public release.

codes. Technical Report, California Univ., Berkeley (USA). Electronics Research Lab.; Lawrence Livermore National Lab., CA (USA), United States, May 1987.

[45] H.D. Rees. Calculation of steady state distribution functions by exploiting stability. *Physics Letters A*, 26(9):416–417, March 1968.

[46] H.D. Rees. Calculation of distribution functions by exploiting the stability of the steady state. *Journal of Physics and Chemistry of Solids*, 30(3):643–655, March 1969.

[47] M. Surendra, D. B. Graves, and G. M. Jellum. Self-consistent model of a direct-current glow discharge: Treatment of fast electrons. *Physical Review A*, 41(2):1112–1125, January 1990.

[48] M. Surendra and D.B. Graves. Particle simulations of radio-frequency glow discharges. *IEEE Transactions on Plasma Science*, 19(2):144–157, April 1991.

[49] Michael E. Jones, Don S. Lemons, Rodney J. Mason, Vincent A. Thomas, and Dan Winske. A Grid-Based Coulomb Collision Model for PIC Codes. *Journal of Computational Physics*, 123(1):169–181, January 1996.

[50] Don S. Lemons, Dan Winske, William Daughton, and Brian Albright. Small-angle Coulomb collision model for particle-in-cell simulations. *Journal of Computational Physics*, 228(5):1391–1403, March 2009.

[51] Wallace M. Manheimer, Martin Lampe, and Glenn Joyce. Langevin Representation of Coulomb Collisions in PIC Simulations. *Journal of Computational Physics*, 138(2):563–584, December 1997.

[52] Bruce I. Cohen, Andris M. Dimits, and David J. Strozzi. A grid-based binary model for coulomb collisions in plasmas. *Journal of Computational Physics*, 234:33–43, February 2013.

[53] F. Pérez, L. Gremillet, A. Decoster, M. Drouin, and E. Lefebvre. Improved modeling of relativistic collisions and collisional ionization in particle-in-cell codes. *Physics of Plasmas*, 19(8):083104, August 2012.

DISTRIBUTION A: Distribution approved for public release.

[54] K. Nanbu and S. Yonemura. Weighted Particles in Coulomb Collision Simulations Based on the Theory of a Cumulative Scattering Angle. *Journal of Computational Physics*, 145(2):639–654, September 1998.

[55] Y. Sentoku and A.J. Kemp. Numerical methods for particle simulations at extreme densities and temperatures: Weighted particles, relativistic collisions and reduced currents. *Journal of Computational Physics*, 227(14):6846–6861, July 2008.

[56] T D Arber, K Bennett, C S Brady, A Lawrence-Douglas, M G Ramsay, N J Sircombe, P Gillies, R G Evans, H Schmitz, A R Bell, and C P Ridgers. Contemporary particle-in-cell approach to laser-plasma modelling. *Plasma Physics and Controlled Fusion*, 57(11):113001, September 2015.

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

[58] Drew Pitney Higginson, Ihor Holod, and Anthony Link. A corrected method for Coulomb scattering in arbitrarily weighted particle-in-cell plasma simulations. *Journal of Computational Physics*, 413:109450, July 2020.

[59] S. Morris, T. Goffrey, K. Bennett, and T. Arber. Improvements to collisional ionization models for particle-in-cell codes. *Physics of Plasmas*, 29(12):123907, December 2022.

[60] 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.

[61] Yong-Ki Kim, José Paulo Santos, and Fernando Parente. Extension of the binary-encounter-dipole model to relativistic incident electrons. *Physical Review A*, 62(5):052710, October 2000.

DISTRIBUTION A: Distribution approved for public release.

[62] J. A. Bearden and A. F. Burr. Reevaluation of X-Ray Atomic Energy Levels. *Reviews of Modern Physics*, 39(1):125–142, January 1967.

[63] Reinhard Bruch, William L. Luken, John C. Culberson, and Kwong T. Chung. K-shell binding energies of B and C. *Physical Review A*, 31(1):503–504, January 1985.

[64] Thomas A. Carlson. *Photoelectron and Auger Spectroscopy*. Springer US, Boston, MA, 1975.

[65] Thomas A. Carlson, C.W. Nestor, Neil Wasserman, and J.D. Mcdowell. Calculated ionization potentials for multiply charged ions. *Atomic Data and Nuclear Data Tables*, 2:63–99, December 1970.

[66] J.P. Desclaux. Relativistic Dirac-Fock expectation values for atoms with $Z = 1$ to $Z = 120$. *Atomic Data and Nuclear Data Tables*, 12(4):311–406, January 1973.

[67] F.P. Larkins. Semiempirical Auger-electron energies for elements $10 \leq Z \leq 100$. *Atomic Data and Nuclear Data Tables*, 20(4):311–387, October 1977.

[68] Wolfgang Lotz. Electron Binding Energies in Free Atoms. *Journal of the Optical Society of America*, 60(2):206, February 1970.

[69] C.C. Lu, T.A. Carlson, F.B. Malik, T.C. Tucker, and C.W. Nestor. Relativistic Hartree-Fock-Slater eigenvalues, radial expectation values, and potentials for atoms, $2 \leq Z \leq 126$. *Atomic Data and Nuclear Data Tables*, 3:1–31, August 1971.

[70] G.C. Rodrigues, P. Indelicato, J.P. Santos, P. Patté, and F. Parente. Systematic calculation of total atomic energies of ground state configurations. *Atomic Data and Nuclear Data Tables*, 86(2):117–233, March 2004.

[71] Kenneth D. Sevier. Atomic electron binding energies. *Atomic Data and Nuclear Data Tables*, 24(4):323–371, October 1979.

[72] D. A. Shirley, R. L. Martin, S. P. Kowalczyk, F. R. McFeely, and L. Ley. Core-electron binding energies of the first thirty elements. *Physical Review B*, 15(2):544–552, January 1977.

DISTRIBUTION A: Distribution approved for public release.

[73] Albert C. Thompson, Janos Kirz, David T. Attwood, Eric M. Gullikson, Malcolm R. Howells, Jeffrey B. Kortright, Yanwei Liu, and Arthur L. Robinson. *X-Ray Data Booklet*. Lawrence Berkeley National Laboratory, 3rd ed edition, September 2009.

[74] Alexander Kramida and Yuri Ralchenko. NIST Atomic Spectra Database, NIST Standard Reference Database 78, 1999.

[75] E. Clementi and D. L. Raimondi. Atomic Screening Constants from SCF Functions. *The Journal of Chemical Physics*, 38(11):2686–2689, June 1963.

[76] E. Clementi, D. L. Raimondi, and W. P. Reinhardt. Atomic Screening Constants from SCF Functions. II. Atoms with 37 to 86 Electrons. *The Journal of Chemical Physics*, 47(4):1300–1307, August 1967.

[77] J. C. Slater. Atomic Shielding Constants. *Physical Review*, 36(1):57–64, July 1930.

[78] M. Guerra, P. Amaro, J.P. Santos, and P. Indelicato. Relativistic calculations of screening parameters and atomic radii of neutral atoms. *Atomic Data and Nuclear Data Tables*, 117–118:439–457, September 2017.

[79] Maria Grazia Pia, Hee Seo, Matej Batic, Marcia Begalli, Chan Hyeong Kim, Lina Quintieri, and Paolo Saracco. Evaluation of Atomic Electron Binding Energies for Monte Carlo Particle Transport. *IEEE Transactions on Nuclear Science*, 58(6):3246–3268, December 2011.

[80] K. L. Bell, H. B. Gilbody, J. G. Hughes, A. E. Kingston, and F. J. Smith. Recommended Data on the Electron Impact Ionization of Light Atoms and Ions. *Journal of Physical and Chemical Reference Data*, 12(4):891–916, October 1983.

[81] H. Deutsch and T.D. Märk. Calculation of absolute electron impact ionization cross-section functions for single ionization of He, Ne, Ar, Kr, Xe, N and F. *International Journal of Mass Spectrometry and Ion Processes*, 79(3):R1–R8, November 1987.

DISTRIBUTION A: Distribution approved for public release.

[82] Robert S. Freund, Robert C. Wetzel, Randy J. Shul, and Todd R. Hayes. Cross-section measurements for electron-impact ionization of atoms. *Physical Review A*, 41(7):3575–3595, April 1990.

[83] C. J. Powell. Cross sections for ionization of inner-shell electrons by electrons. *Reviews of Modern Physics*, 48(1):33–47, January 1976.

[84] H. Tawara and T. Kato. Total and partial ionization cross sections of atoms and ions by electron impact. *Atomic Data and Nuclear Data Tables*, 36(2):167–353, March 1987.

DISTRIBUTION A: Distribution approved for public release.

Project 2: Development of a new forward adaptively refined and regularized semi-Lagrangian integral/Green's function method for solving the Vlasov-Poisson system accelerated with a hierarchical treecode

DISTRIBUTION A: Distribution approved for public release.

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

The FARSIGHT Vlasov-Poisson code

Ryan Sandberg[a], Robert Krasny[a], Alexander G. R. Thomas[b]

[a]*Department of Mathematics, University of Michigan, 500 S. State Street, Ann Arbor 48109, USA*
[b]*Department of Nuclear Engineering and Radiological Sciences, University of Michigan, 500 S. State Street, Ann Arbor 48109, USA*

ARTICLE INFO

Article history:

ABSTRACT

A forward semi-Lagrangian scheme called FARSIGHT is presented for modeling collisonless 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 which are tracked by 4th order Runge-Kutta time-stepping. 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. Numerical results are presented for Landau damping, two-stream instability, and halo formation for a mismatched thermal equilibrium sheet beam in a focusing channel.

© 2023 Elsevier Inc. All rights reserved.

1. Introduction

Collisionless plasmas modeled using the 1D1V Vlasov-Poisson system of equations arise in many applications including shock acceleration [1, 2], beam physics [3, 4], inertial confinement fusion [5, 6], and plasma-wakefield acceleration [7, 8]. Here we consider the kinetic model of a one-component electron plasma described by the 1D1V Vlasov-Poisson equations in non-dimensional form,

$$\partial_t f + v\partial_x f + F\partial_v f = 0, \quad F = qE, \quad E = -\partial_x \phi, \quad -\partial_x^2 \phi = \rho, \tag{1}$$

where $f(x, v, t)$ is the number density of electrons in (x, v) phase space, $F(x, t)$ is the force on an electron, $q = -1$ is the electron charge, $E(x, t)$ is the electric field, $\phi(x, t)$ is the electric potential, and $\rho(x, t)$ is the total charge density. Periodic boundary conditions are enforced on $0 \le x \le L$. The electron number density is normalized,

$$\int_{-\infty}^{\infty} \int_0^L f(x, v, t) dx dv = 1, \tag{2}$$

e-mail: ryansand@umich.edu (Ryan Sandberg)

DISTRIBUTION A: Distribution approved for public release.

and the charge density is given by

$$\rho(x,t) = q\int_{-\infty}^{\infty} f(x,v,t)dv + \frac{1}{L}, \tag{3}$$

where the first term corresponds to the electrons and the second term is a static uniform distribution of positive ions to enforce neutrality,

$$\int_0^L \rho(x,t)dx = 0. \tag{4}$$

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.

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 [9, 10]. PIC is in widespread use, but difficulties may arise due to discrete particle noise unless sufficiently many particles are used [4, 11]. In addition, particle-grid interactions may induce thermal heating and energy drift [9, 10]. 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 [12], 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 [13]. For electrostatic 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 [14, 15, 16].

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 [5, 4] and the dynamics of electron holes [17, 18]. Eulerian approaches include finite difference methods [19], finite volume methods [20, 21], finite element and discontinuous Galerkin methods [22], and spectral methods [23, 24, 8, 25].

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 [26, 27, 28]. In these codes the spatial discretization is often done using discontinuous Galerkin methods [29, 30, 31] or weighted essentially non-oscillatory (WENO) methods [32, 33, 34]. While early semi-Lagrangian methods used operator splitting [27], multidimensional methods are employed recently [35, 36]. 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 [37, 38], multi-grid remeshing [39], block decomposition [40] and multiresolution wavelet representations [41, 42, 43].

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 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 [44].

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 [45]; 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

DISTRIBUTION A: Distribution approved for public release.

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 $O(N^2)$ operations using direct summation, but the cost is reduced to $O(N \log N)$ 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 article is as follows. Section 2 discusses several preliminary topics including the flow map, electric field integral, kernel regularization, and the electric field following a particle. Section 3 develops the numerical method. Section 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 5 summarizes the work.

2. Derivation of regularized electric field integral

This section presents the derivation of the regularized electric field kernel needed for the numerical method. This encompasses the flow map, electric field integral, kernel regularization, and electric field at a particle position.

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

$$\partial_t x(x_0, v_0, t) = v(x_0, v_0, t), \quad \partial_t v(x_0, v_0, t) = qE(x(x_0, v_0, t), t), \tag{5}$$

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

$$f(x(x_0, v_0, t), v(x_0, v_0, t), t) = f_0(x_0, v_0), \tag{6}$$

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 $\mathcal{E}$,

$$Q = \int_0^L \int_{-\infty}^{\infty} f dv dx, \quad I = \int_0^L \int_{-\infty}^{\infty} v f dv dx, \quad \mathcal{E} = \frac{1}{2}\int_0^L \int_{-\infty}^{\infty} v^2 f dv dx + \frac{1}{2}\int_0^L E^2 dx. \tag{7}$$

2.2. Integral solution of Poisson equation

This section derives an integral expression for the potential $\phi(x, t)$ satisfying the Poisson equation in Eq. (1) with periodic boundary conditions on $0 \le x \le L$ and charge density in Eq. (3) satisfying charge neutrality. Consider the free-space Green's function for the 1D Laplace operator,

$$g(x, y) = -\frac{1}{2}|x - y|. \tag{8}$$

The potential is the sum of two terms, a particular solution ϕ_p given by convolving the Green's function with the charge density, and a homogeneous solution ϕ_0,

$$\phi = \phi_p + \phi_h, \quad \phi_p(x, t) = \int_0^L g(x, y)\rho(y, t)dy, \quad \phi_h(x, t) = ax + b.$$

Since the potential is determined up to an arbitrary additive constant, we set $b = 0$. Then charge neutrality ensures that the potential gradient is periodic,

$$\partial_x \phi(L, t) - \partial_x \phi(0, t) = \int_0^L \partial_{xx}\phi(x, t)dx = \int_0^L -\rho(x, t)dx = 0, \tag{9}$$

and enforcing periodicity of the potential yields an expression for the remaining constant,

$$a = \frac{1}{L}(\phi_p(0, t) - \phi_p(L, t)) = \frac{1}{L}\int_0^L (g(0, y) - g(L, y))\rho(y, t)dy = \int_0^L \left(\frac{1}{2} - \frac{y}{L}\right)\rho(y, t)dy = -\frac{1}{L}\int_0^L y\rho(y, t)dy. \tag{10}$$

The integral form for the potential is,

$$\phi(x, t) = \int_0^L \left(g(x, y) - \frac{xy}{L}\right)\rho(y, t)dy. \tag{11}$$

DISTRIBUTION A: Distribution approved for public release.

2.3. Electric field integral

The electric field is

$$\begin{aligned} E(x,t) = -\phi_x(x,t) &= \int_0^L \Big(-g_x(x,y) + \frac{y}{L}\Big)\rho(y,t)dy \\ &= \int_0^L \Big(-g_x(x,y) + \frac{y}{L}\Big)\Big(q\int_{-\infty}^{\infty} f(y,v,t)dv + \frac{1}{L}\Big)dy \\ &= q\int_0^L\int_{-\infty}^{\infty}\Big(-g_x(x,y) + \frac{y}{L}\Big)f(y,v,t)dvdy + \int_0^L\Big(-g_x(x,y) + \frac{y}{L}\Big)dy\frac{1}{L} \\ &= q\int_0^L\int_{-\infty}^{\infty}\Big(-g_x(x,y) + \frac{y}{L}\Big)f(y,v,t)dvdy + \frac{x}{L} \\ &= \int_{\infty}^{\infty}\int_0^L\Big(g_x(x,y) + \frac{x-y}{L}\Big)f(y,v,t)dydv, \\ &= \int_{\infty}^{\infty}\int_0^L K(x,y)f(y,v,t)dydv \end{aligned} \tag{12}$$

where the electric field kernel is

$$K(x,y) = g_x(x,y) + \frac{x-y}{L}, \quad g_x(x,y) = -\frac{1}{2}\text{sign}(x-y). \tag{13}$$

To regularize K while ensuring that the regularization respects periodicity, we observe that $K(x,y)$ can be written as

$$K(x,y) = k\Big(\frac{x-y}{L}\Big) = -\frac{1}{2}\text{sign}\Big(\frac{x-y}{L}\Big) + \frac{x-y}{L}. \tag{14}$$

The function $k(s)$ has a domain $-1 \le s \le 1$,, but examining the function we see that it has a 1-periodic extension, ($k(s+1) = k(s)$), piecewise linear with slope 1 between the integers $s = 0, \pm1, \pm2, \ldots$ 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. 1f).

2.4. 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. (5); 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 [46, 47, 48]. We start by writing the elementary kernel as the sum of two terms,

$$k(s) = k_1(s) + k_2(s), \quad k_1(s) = -\frac{1}{2}\text{sign}(s), \quad k_2(s) = s, \quad -\frac{1}{2} \le s \le \frac{1}{2}, \tag{15}$$

and extend periodically. That is, $k_1(s)$ is the tophat function shown in Fig. 1b and $k_2(s) = s$ extended periodically is the shifted sawtooth shown in Fig. 1b. Note that $k_1(s)$ 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. 1d), while $k_2(s)$ can be understood as the field due to a periodic array of negative discrete charges at the half-integers with a uniform neutralizing positive background (Fig. 1e). Based on this splitting we define the regularized elementary kernel,

$$k_\epsilon(s) = k_{1,\epsilon}(s) + k_2(s), \quad k_{1,\epsilon}(s) = -\frac{c_\epsilon}{2}\frac{k_2(s)}{(k_2(s)^2 + \epsilon^2/L^2)^{1/2}}, \quad c_\epsilon = (1 + 4\epsilon^2/L^2)^{1/2}, \tag{16}$$

where $\epsilon > 0$ is the regularization parameter. Figure 1i depicts $k_\epsilon(s)$ for a representative value of ϵ, and Figs. 1g,h show the two components, where in Fig. 1g $k_{1,\epsilon}(s)$ is a regularized version of the tophat and in Fig. 1h $k_2(s)$ has not changed. Figure 1l shows the charge density $k'_\epsilon(x)$ of the regularized elementary kernel, which is the superposition of the charge densities $k'_{1,\epsilon}(x)$ and $k'_2(x)$ in Figs. 1j,k.

DISTRIBUTION A: Distribution approved for public release.

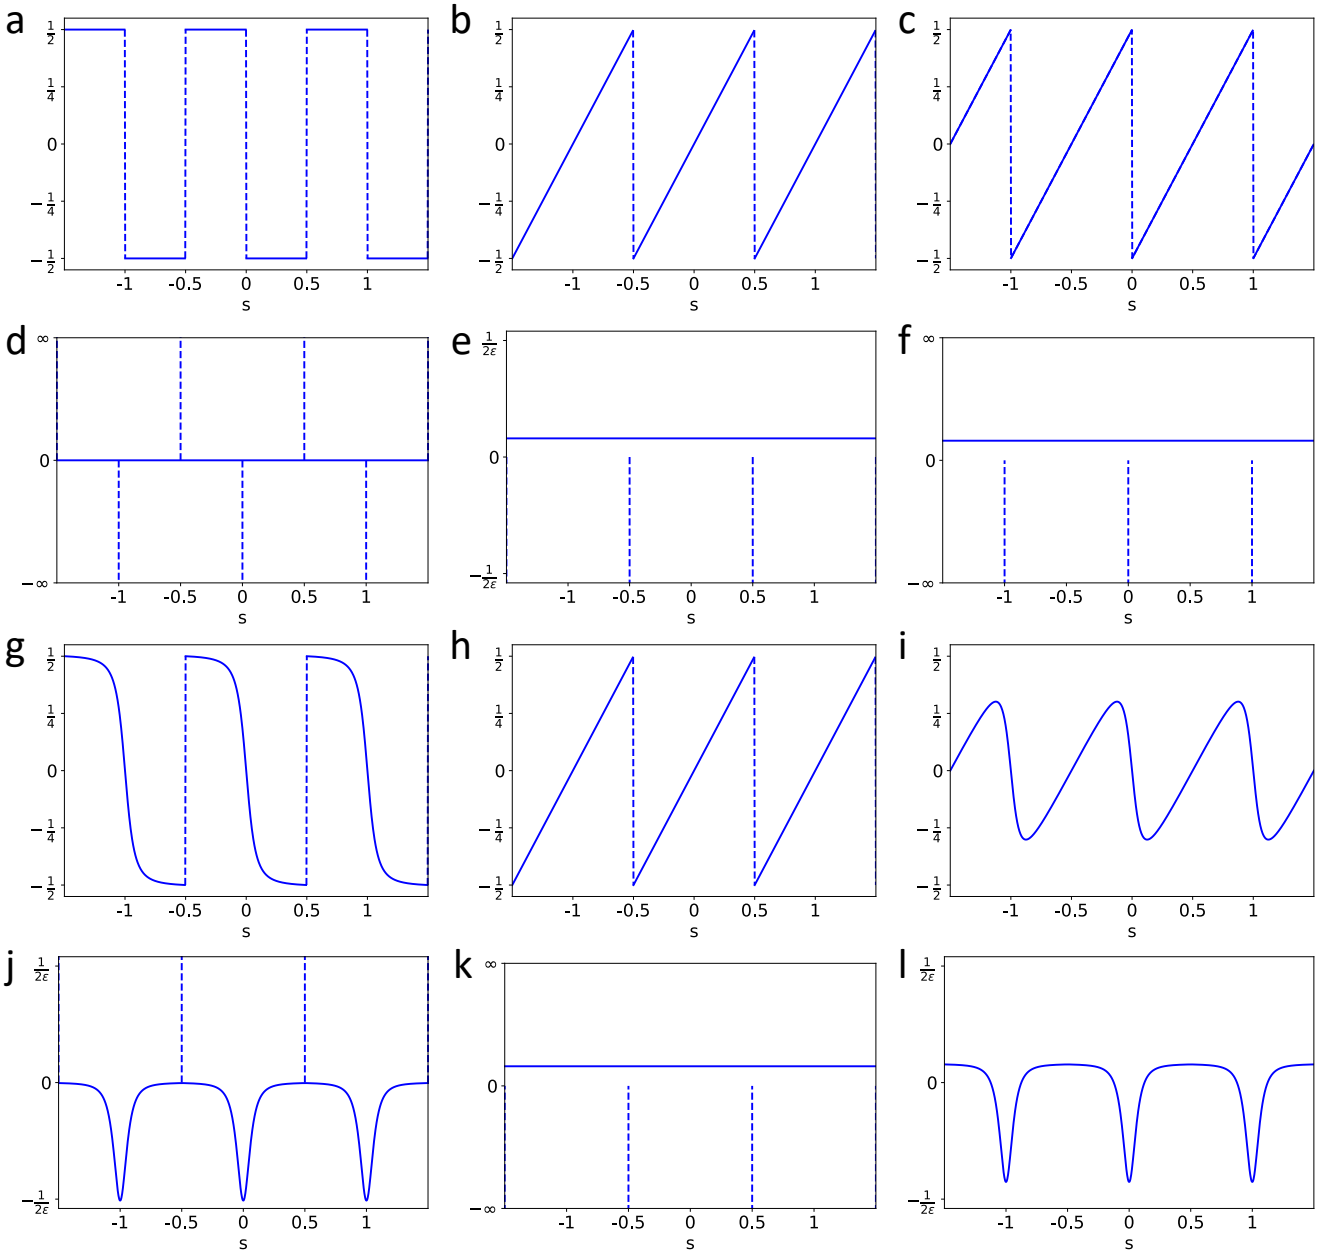


Fig. 1: Elementary electric field kernel $k(s) = k_1(s) + k_2(s)$ (a) tophat $k_1(s)$, (b) shifted sawtooth $k_2(s)$, (c) sawtooth $k(s)$ in Eq. (**??**), splitting $k(s) = k_1(s) + k_2(s)$ in Eq. (15) associated charge density (d) $k_1'(s)$, (e) $k_2'(s)$, (f) $k'(s)$, regularized elementary electric field kernel (g) regularized tophat $k_1^\epsilon(s)$, (h) shifted sawtooth $k_2(s)$, (i) $k_\epsilon(s)$ in Eq. (17), associated charge density (j) $k_{1,\epsilon}'(s)$, (k) $k_2'(s)$, (l) $k_\epsilon'(s)$.

In principle, other forms of regularization are possible, but the form given here was found to be convenient. Note that $k_\epsilon(s) \to k(s)$ as $\epsilon \to 0$ 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 2nd derivative $k_\epsilon''(s)$ of order $O(\epsilon^2)$ 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 0 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

$$K_\epsilon(x, y) = k_\epsilon\left(\frac{x - y}{L}\right), \tag{17}$$

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

DISTRIBUTION A: Distribution approved for public release.

2.5. Electric field at a particle position

The electric field at a particle position is

$$E(x(x_0, v_0, t), t) = \int_0^L \int_{-\infty}^{\infty} K(x(x_0, v_0, t), y) f(y, v, t) dv dy. \tag{18}$$

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

$$E(x(x_0, v_0, t), t) = \int_0^L \int_{-\infty}^{\infty} K(x(x_0, v_0, t), x(y_0, u_0, t)) f(x(y_0, u_0, t), v(y_0, u_0, t), t) J(y_0, u_0, t) du_0 dy_0 \ . \tag{19}$$

Using Eq. (6) 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

$$E(x(x_0, v_0, t), t) = \int_0^L \int_{-\infty}^{\infty} K(x(x_0, v_0, t), x(y_0, u_0, t)) f_0(y_0, u_0) du_0 dy_0 \ . \tag{20}$$

Hence Eq. (20) expresses the electric field as a convolution of the electric field kernel with the initial electron distribution.

3. Numerical Method

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

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. 2a. The panels have a tree structure and, as in previous work in fluid dynamics [45], are adaptively refined; for example in Fig. 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,

$$\text{var}(P) = \max_{i,j=1:3} f_{ij} - \min_{i,j=1:3} f_{ij} \le \epsilon_{\text{var}}, \tag{21}$$

where ϵ_{var} is a user-specified tolerance. If the criterion in Eq. (21) 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. (21) is known analytically, while at later times when remeshing is applied it is computed by interpolation as explained below.

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 1.

DISTRIBUTION A: Distribution approved for public release.

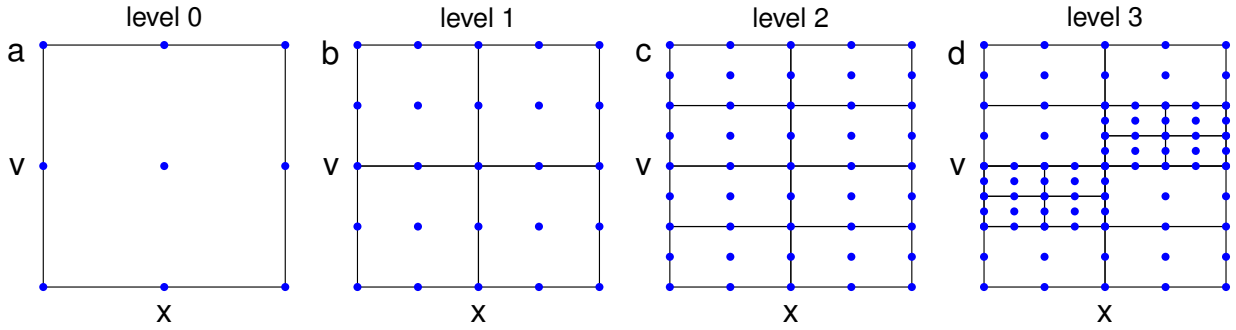


Fig. 2: Phase space discretization by Lagrangian particles and panels, each panel is a rectangle with a 3×3 particle grid (•), panels have a tree structure and are adaptively refined, (a) level 0, root panel comprising the entire computational domain, (b) level 1, uniformly refined in x and v, (c) level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

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

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. (20) 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. The trapezoid rule is spectrally accurate in the case of periodic functions and so higher order discretizations were not studied in the initial FARSIGHT development. When substituted into Eq. (5), this yields a set of ODEs for the motion of the particles in phase space,

$$x_i' = v_i, \quad v_i' = q \sum_{j=1}^{N} K_\epsilon(x_i, x_j) f_{0j} \Delta x_{0j} \Delta v_{0j} w_j, \quad x_i(0) = x_{0i}, \quad v_i(0) = v_{0i}, \quad i = 1 : N, \tag{22}$$

where $K_\epsilon(x, v)$ is the regularized kernel in Eq. (17), $f_{0j} = f_0(x_{0j}, v_{0j})$ is the electron density of a particle, $\Delta x_{0j}, \Delta v_{0j}$ are panel dimensions, w_j are the trapezoid rule weights (1/16 for panel vertices, 1/8 for edge midpoints, 1/4 for the panel midpoint), and the minus sign in front of the sum is from the Lorenz expression for the force. In this work the ordinary differential equations (ODEs) are solved using the classical 4th order Runge-Kutta 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.

Discrete analogues of the conserved quantities are obtained by applying the trapezoid rule to the integrals in Eq. (7),

$$Q_N = \sum_{i=1}^{N} f_i \Delta x_i \Delta v_i w_i, \quad I_N = \sum_{i=1}^{N} v f_i \Delta x_i \Delta v_i w_i, \quad \mathcal{E}_N = \frac{1}{2} \sum_{i=1}^{N} v^2 f_i \Delta x_i \Delta v_i w_i + \frac{1}{2} \sum_{\ell=1}^{M} E_\ell^2 \Delta x_\ell w_\ell, \tag{23}$$

where the sums for Q_N, I_N, and the first term of $\mathcal{E}_N$ arise from integrals over phase-space, the sum for the second term of $\mathcal{E}_N$ arises from an integral over physical space, and w_i, w_ℓ are the Trapezoid rule quadrature weights in each case.

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.

DISTRIBUTION A: Distribution approved for public release.

Figure 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. (22), 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 [39] and also occurs in vortex methods for incompressible fluid flow [49, 48].

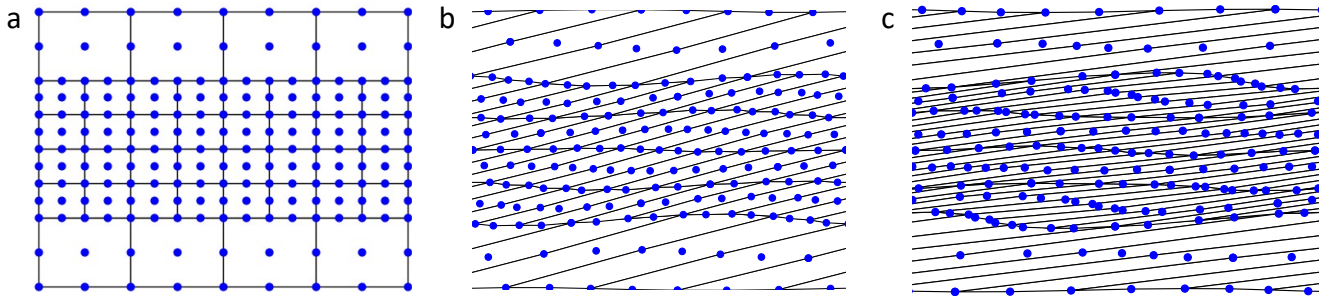


Fig. 3: Landau damping simulation by Eq. (22) illustrating how a fixed set of Lagrangian particles and panels is progressively deformed, (a) $t = 0$, (b) $t = 3$, (c) $t = 6$.

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 [50] as well as PIC plasma simulations [39, 51].

3.4. Remeshing

Figure 4 illustrates the remeshing scheme. For simplicity, assume the particles initially lie on a uniform mesh with regular panels as in Fig. 4a (in practice they can lie on an adaptively refined mesh as in Fig. 2d or Fig. 3a). After some time, the particles and panels are deformed as in Fig. 4b. The remeshing scheme replaces these particles and panels by new particles and panels on a new uniform regular mesh, as in Fig. 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.

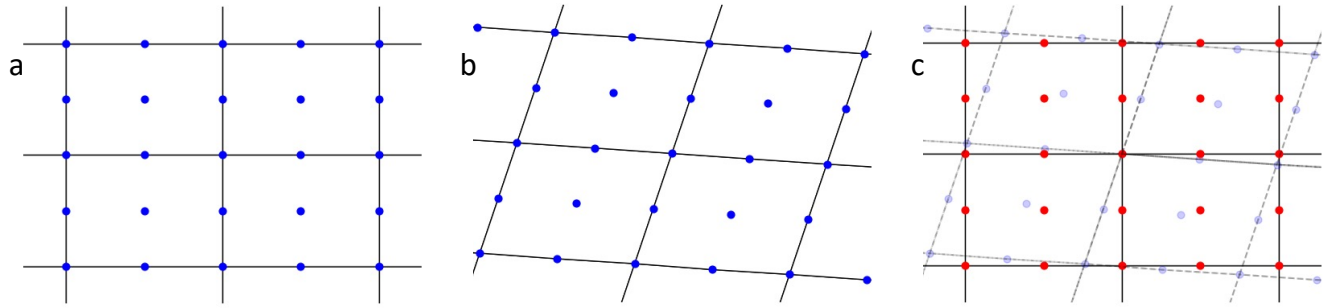


Fig. 4: Remeshing, (a) regular particles/panels at initial time, (b) irregular particles/panels after evolving for some time, (c) new regular particles/panels (red), old irregular particles/panels (light blue), electron distribution is transferred from old particles to new particles using biquadratic 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

DISTRIBUTION A: Distribution approved for public release.

employs neighbor panel information from the previous new particle. In problems with periodic boundary conditions, the neighbor panel search takes x-periodicity into account.

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,

$$p(x,v) = \sum_{i,j=1}^{3} c_{ij}(x - x_{22}^{old})^{i-1}(v - v_{22}^{old})^{j-1}, \tag{24}$$

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

$$p(x_{ij}^{old}, v_{ij}^{old}) = f_{ij}^{old}, \quad i, j = 1:3. \tag{25}$$

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 [29, 39, 51]. 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.

3.6. Barycentric Lagrange treecode

The discrete approximation of the electric field in Eq. (22) has the form of an N-body interaction,

$$E_i = \sum_{j=1}^{N} K_\epsilon(x_i, x_j) q_j, \quad q_j = q f_{0j} \Delta x_{0j} \Delta v_{0j} w_j, \quad i = 1:N, \tag{26}$$

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 at each remeshing. If the field E_i is evaluated by direct summation using loops over i, j, then the operation count is $O(N^2)$, which is costly for large N. In this work we employ the barycentric Lagrange treecode (BLTC) to reduce the operation count to $O(N \log N)$ [52, 53]. More particularly, we use the implementation developed in [44].

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. 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 [54]. The resulting approximation has the form,

$$E_i \approx \sum_{C_{\text{near}}} \sum_{x_j \in C_{\text{near}}} K_\epsilon(x_i, x_j) q_j + \sum_{C_{\text{far}}} \sum_{k=0}^{n} K_\epsilon(x_i, s_k) \hat{q}_k, \quad \hat{q}_k = \sum_{x_j \in C_{\text{far}}} L_k(x_j) q_j, \tag{27}$$

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 5b,c depicts the particle-particle interactions.

DISTRIBUTION A: Distribution approved for public release.

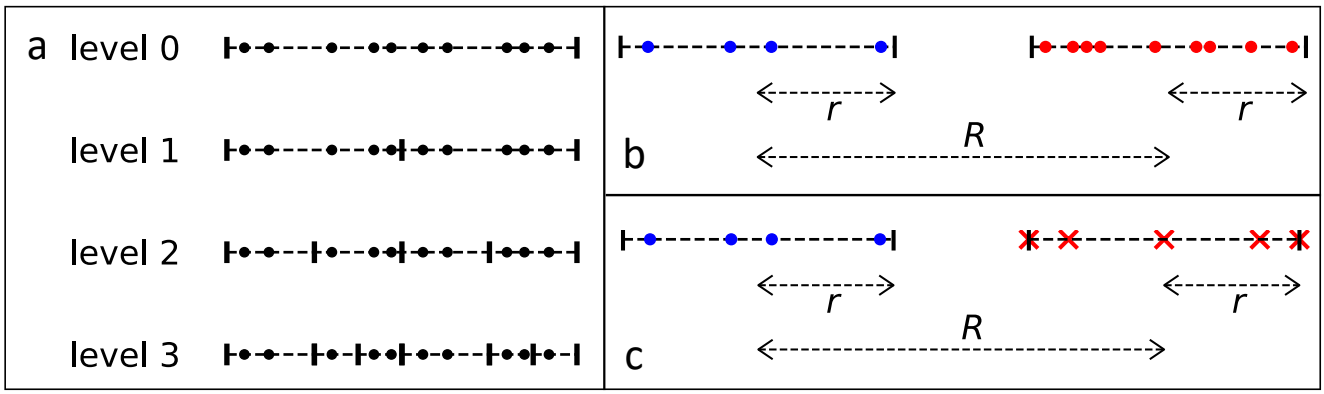


Fig. 5: Treecode schematic, (a) the particles are partitioned into clusters (subintervals of $[0, L]$) with a hierarchical tree structure, in (b,c) the target particles (left,•) interact with (b) source particles (right,•) if there are few enough source particles or with (c) Chebyshev points (right,×) if there are enough source particles (right, cluster radius r, particle-cluster distance R.

The barycentric form of the Lagrange polynomials is employed,

$$L_k(x) = \frac{\dfrac{w_k}{x - s_k}}{\displaystyle\sum_{k'=0}^{n} \frac{w_{k'}}{x - s_{k'}}}, \quad w_k = (-1)^k \delta_k, \quad \delta_k = \begin{cases} 1/2, & k = 0, n, \\ 1, & k = 1 : n-1, \end{cases} \tag{28}$$

which is computationally stable and efficient [54]. Note that the proxy charges $\hat{q}_k$ of a cluster C in Eq. (27) 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.

A small issue occurs if too many particles have the same x-coordinate, which would interfere with the definition of the leaf clusters in the tree. In practice this only happens in the first stage of the Runge-Kutta timestepping. To avoid any problem, particles having the same x-coordinate are lumped together into a single particle with the appropriate weight.

3.7. GPU implementation

A key feature of the BLTC is that the electric field approximation E_i in Eq. (27) 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 Open ACC [44]. 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.

3.8. BLTC performance

This section documents the accuracy and efficiency of the BLTC 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. (26), with particles distributed on a cold wave,

$$x_i = i\Delta x + 0.5 \cos(2\pi i \Delta x), \quad \Delta x = 1/N, \quad f_i = 1, \quad i = 1 : N. \tag{29}$$

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 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

DISTRIBUTION A: Distribution approved for public release.

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.

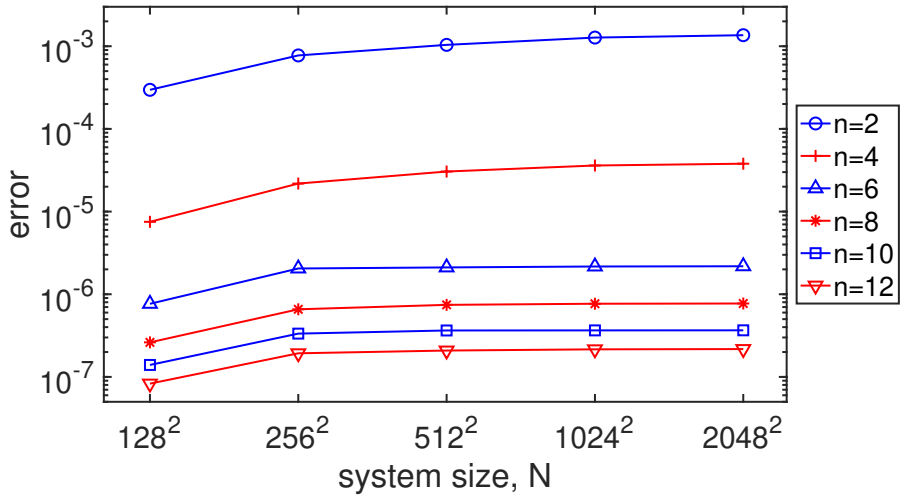


Fig. 6: BLTC performance, single field evaluation, particles given by Eq. (29), error versus system size N, MAC $\theta = 0.6$, degree $n = 2 : 2 : 12$.

Figure 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 7a shows results for 36 CPU cores and Fig. 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 paper.

Figure 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 MACX $\theta = 0.6$ 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.

4. Numerical Results

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

$$f_M(v; v_{\text{th}}) = \frac{1}{\sqrt{2\pi} v_{\text{th}}} \exp\left(-\frac{v^2}{2v_{\text{th}}^2}\right), \quad -\infty < v < \infty, \tag{30}$$

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.

DISTRIBUTION A: Distribution approved for public release.

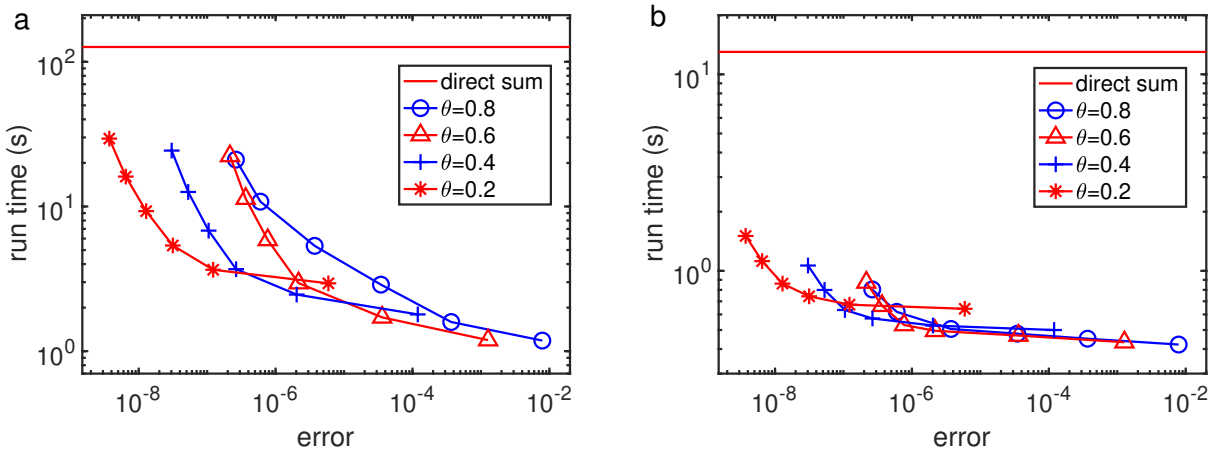


Fig. 7: BLTC performance, single field evaluation, particles given by Eq. (29), run time (s) versus error, $N = 1024^2$ particles, MAC $\theta = 0.2, 0.4, 0.6, 0.8$ constant on each curve, symbols give degree $n = 2, 4, 6, 8, 10, 12$ increasing from right to left on each curve, (a) 36 CPU cores, (b) one GPU, note different vertical scales in (a) and (b), horizontal lines give direct sum run time (s).

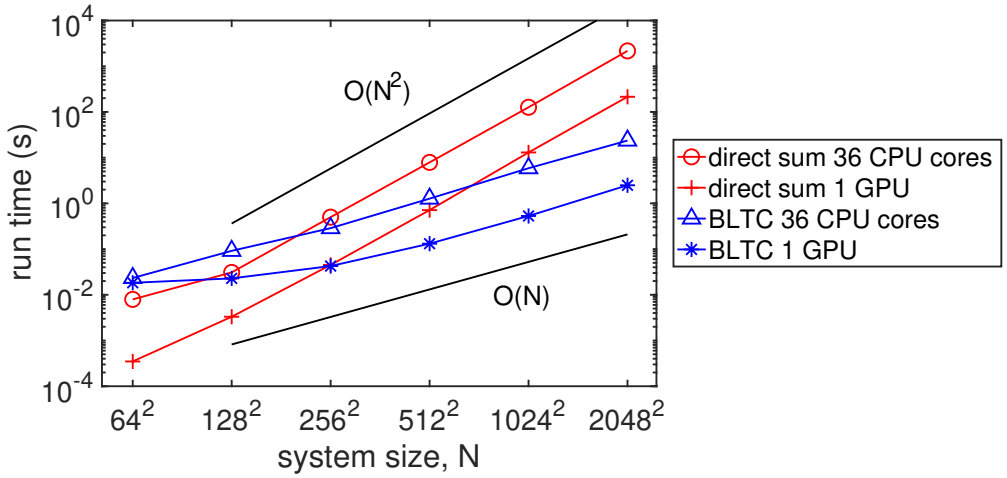


Fig. 8: BLTC performance, single field evaluation, particles given by Eq. (29), MAC $\theta = 0.6$, degree $n = 8$ yielding error ≈1e-6, run time (s) vs. system size N, direct sum and BLTC on 36 CPU cores and 1 GPU.

These problems are posed with periodic boundary conditions in x-space and cutoff $v_{\max} = 6$ 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.

4.1. Landau damping

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

$$f_0(x, v) = f_M(v; v_{\text{th}})(1 + \alpha \cos(kx)), \qquad (x, v) \in [0, 4\pi] \times [-6, 6], \tag{31}$$

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

4.1.1. Weak Landau damping

First we consider small perturbation amplitude, $\alpha = 0.01$. Figure 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 [55],

and the computed results are in good agreement with the predicted decay rate $\gamma = 0.1533$ [29] indicated by the dashed line.

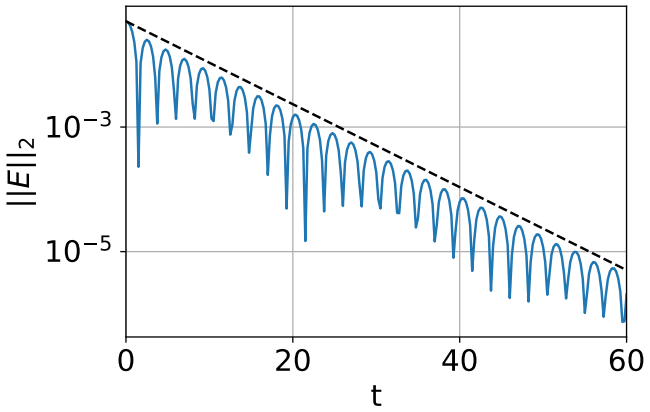


Fig. 9: Weak Landau damping, 2-norm of electric field versus time, phase space discretization $N = 256^2$, time step $\delta t = 0.25$, regularization parameter $\epsilon = 0.1$, slope of the dashed line is the benchmark decay rate $\gamma = 0.1533$ [29], GPU run time 34.6 s.

Figure 10 shows the dependence of the electric field decay on the numerical parameters compared with the theoretical decay rate indicated by the black dashed line in each figure. 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 indicated by the dashed black line as the numerical parameters are refined.

Next we consider the electron distribution. Figure 11a shows the initial distribution, where the perturbation is too small to be seen. Figures 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 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 [56].

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 each case, one parameter was varied and the other two were kept fixed. Figure 13a shows that the error scales like $O(\Delta x^3)$, the expected rate for biquadratic interpolation and 2D trapezoid quadrature. Figure 13b shows that the error scales like $O(\Delta t^4)$, the expected rate for 4th order Runge-Kutta time-stepping. Figure 13c shows that the error scales like $O(\epsilon^2)$, which is the rate at which the regularized electric field kernel in Eq. (16) converges away from the singularity. Convergence with respect to a regularization parameter has been considered in analyses of particle methods [57, 48, 58] (Beale, Majda, Cottet, Raviart), but we are not aware of any rigorous result for the present case.

To end the discussion of weak Landau damping, we consider the conserved quantities in Eqs. (23). Figure 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. 14 the errors in (a) total charge $|Q_N - Q_{N0}|$, (b) total current $|I_N - I_{N0}|$, and (c) total energy $|\mathcal{E}_N - \mathcal{E}_{N0}|$ are shown. The error in the conserved quantities is limited by the treecode error, as the treecode parameters are not varied here. To demonstrate further control, the treecode parameters would also need to be refined.

DISTRIBUTION A: Distribution approved for public release.

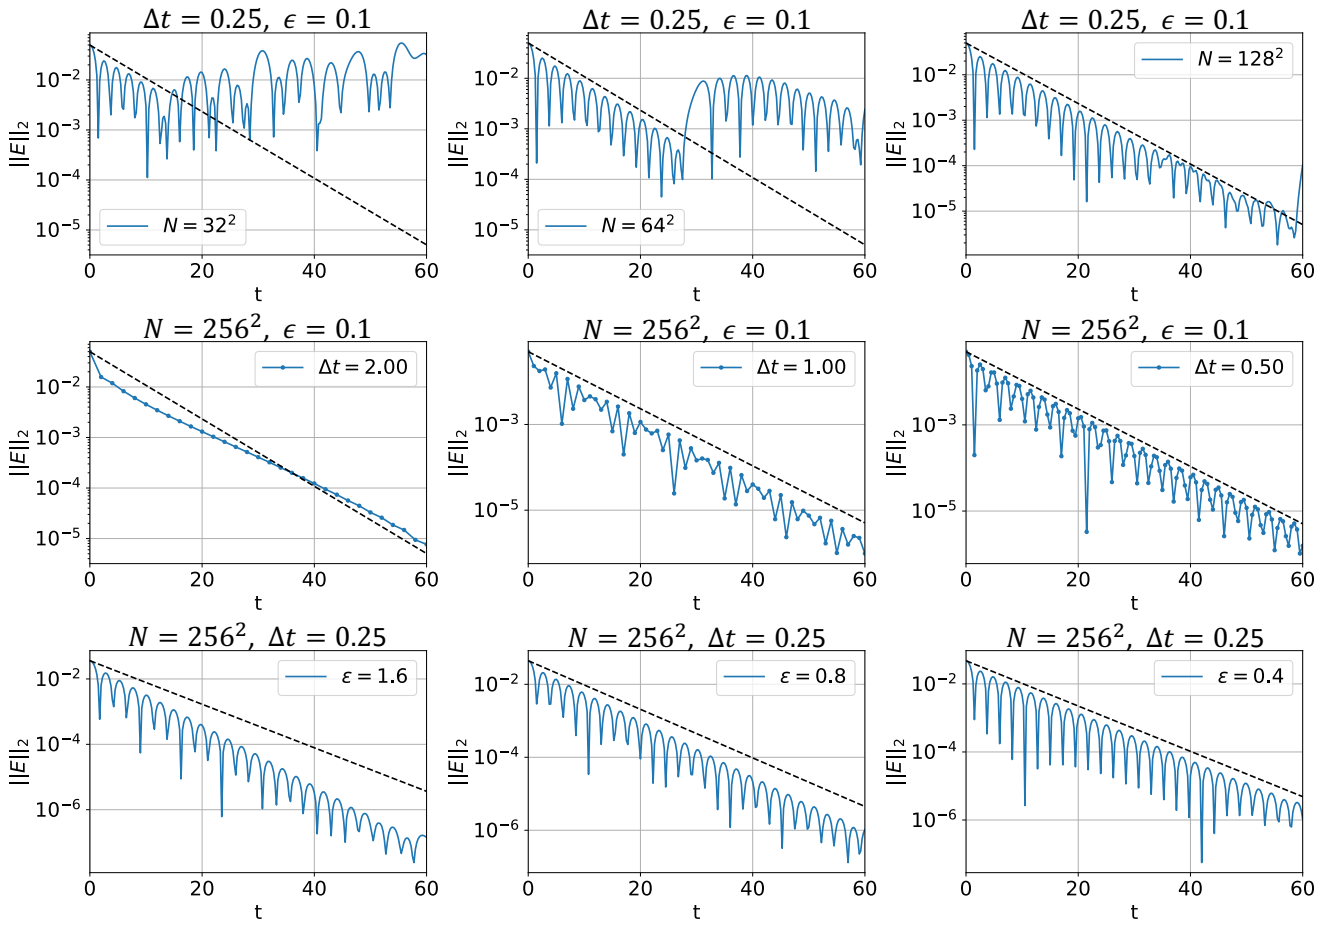


Fig. 10: Weak Landau damping, 2-norm of electric field versus time, effect of numerical parameters on decay rate, theoretical decay rate indicated by dashed black line, top row: phase space discretization $N = 32^2, 64^2, 128^2$, middle row: time step $\Delta t = 2, 1, 0.5$, bottom row: regularization parameter $\epsilon = 1.6, 0.8, 0.4$.

4.1.2. Strong Landau damping

Figure 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 documented in the literature [59, 60, 33].

Figure 16 shows the corresponding electron distribution, which exhibits significant nonlinear 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 17 shows the behavior of conservation metrics with refinement in parameters. Fig. 17(a) shows the variation in total charge $|Q_N - Q_{N0}|$, (b) total current $|I_N - I_{N0}|$, and (c) total energy $|\mathcal{E}_N - \mathcal{E}_{N0}|$. We note that all errors are controlled as numerical parameters are refined. The conservation in energy is comparable with the literature [29, 61, 62].

The calculations presented so far used a uniform mesh. In Fig. 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 the highest level are too small to be seen clearly in the figure, so instead they are shaded by level of refinement. In Fig. 19 a smaller region of phase-space at time $t = 60$ is shown so that the panels can be seen. Fig. 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. 19(c) shows the actual panel outlines as opposed to a heat map coloring panels by AMR level. This affords a

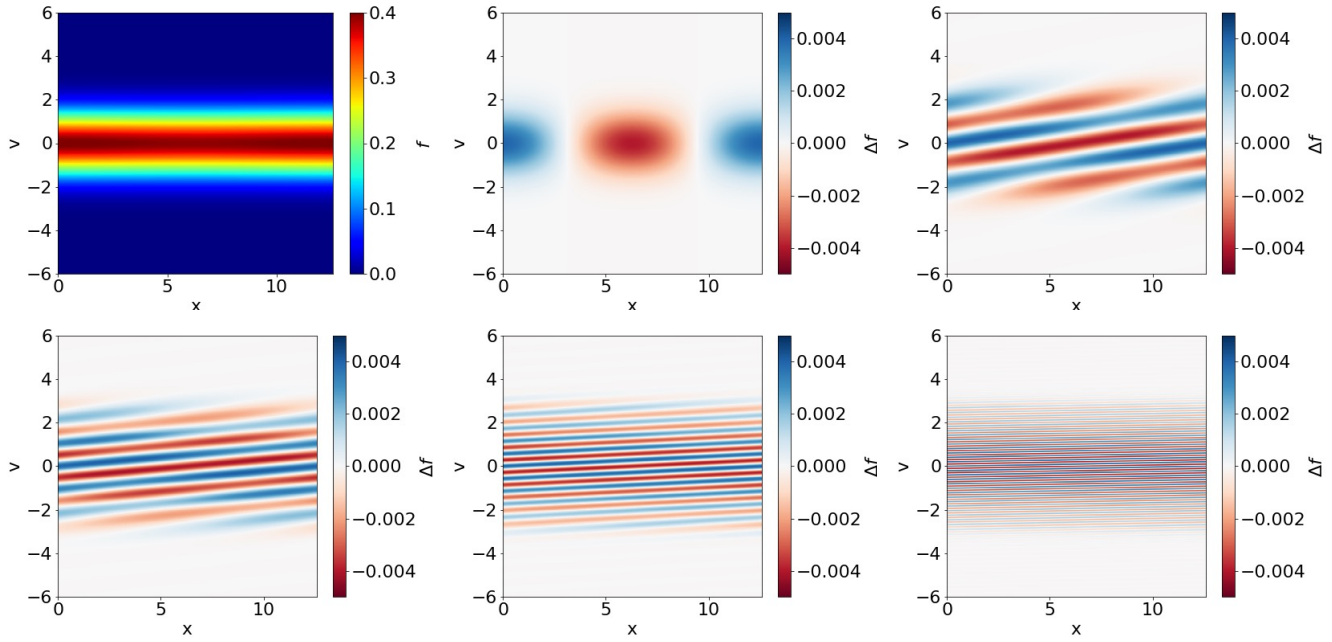


Fig. 11: Weak Landau damping, (a) initial electron distribution $f_0(x, v)$, (b)-(f) perturbation from Maxwellian, $\delta f(x, v, t) = f(x, v, t) - f_M(v)$, time $t \in [0, 60]$, phase space discretization $N = 1024^2$, time step $\Delta t = 0.125$, regularization parameter $\epsilon = 0.05$, run time 1372 s.

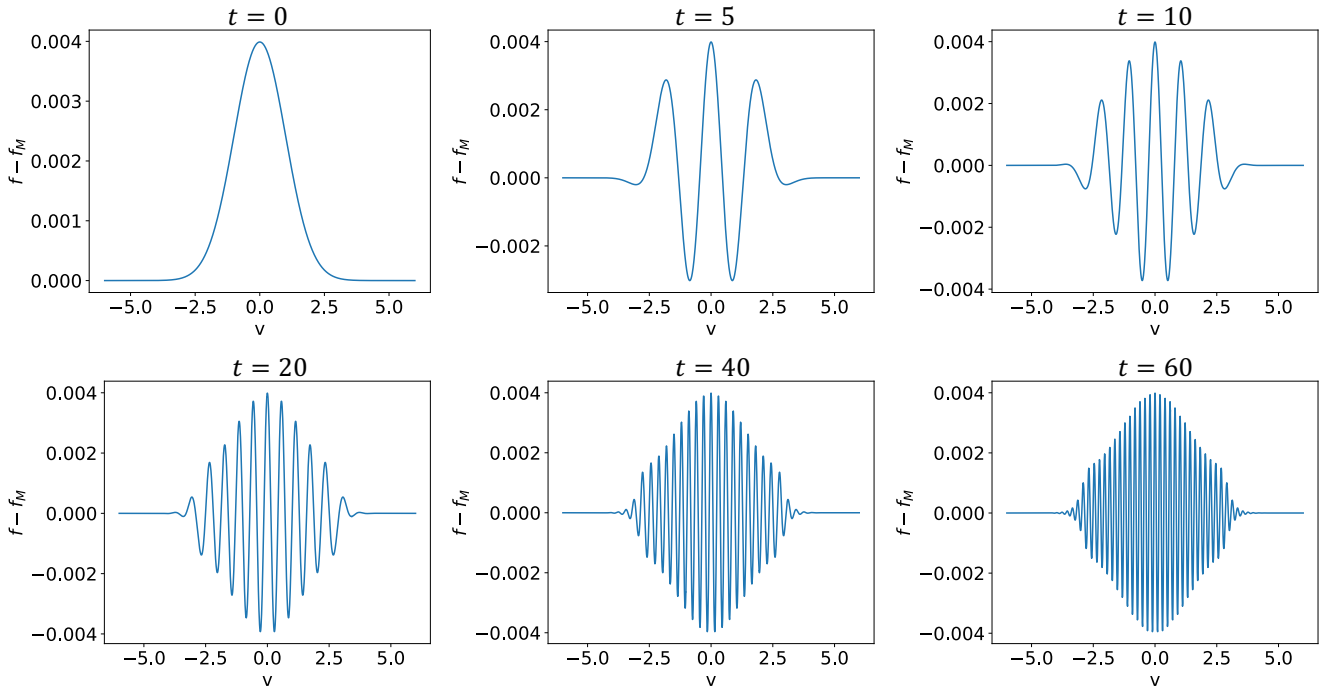


Fig. 12: Weak Landau damping, slices through electron density deviation from Maxwellian, $\delta f(0, v, t) = f(0, v, t) - f_M(v)$, time $t \in [0, 60]$, same parameters as Fig. 11.

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. 16, and the AMR calculation performed with $N = 16 \times 64 + 5$ levels of refinement, shown in Fig. 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× faster than the uniform

DISTRIBUTION A: Distribution approved for public release.

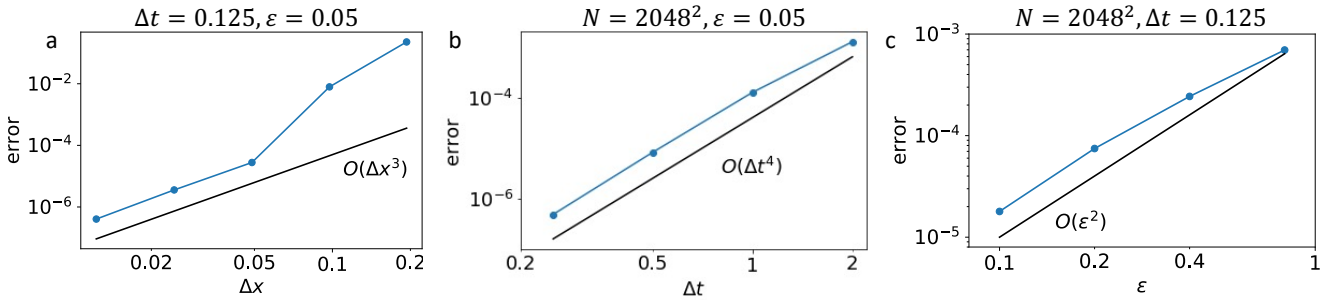


Fig. 13: Weak Landau damping, convergence of electron density with respect to numerical parameters, time $t = 60$, reference solution computed with phase space discretization $N = 2048^2$, time step $\Delta t = 0.125$, regularization parameter $\epsilon = 0.05$, solid line shows asymptotic error scaling, (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 = $O(\epsilon^2)$.

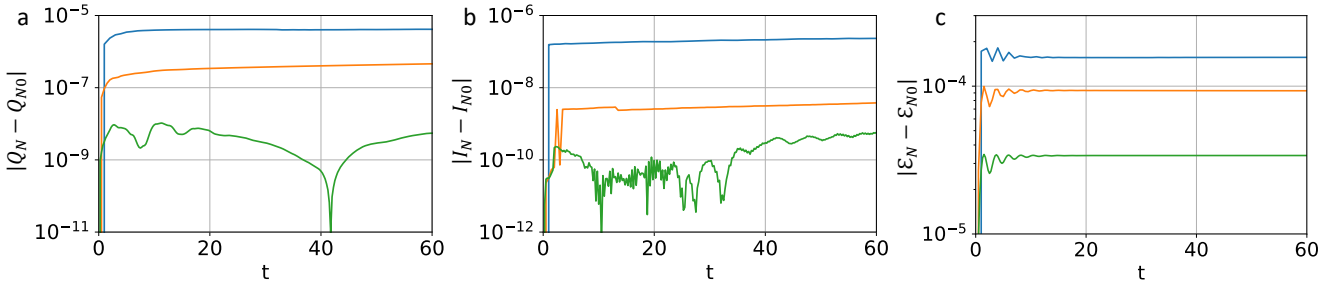


Fig. 14: Weak Landau damping, errors in conserved quantities vs. time, effect of numerical parameters, (a) total charge error $|Q_N - Q_{N0}|$, (b)total current error $|I_N - I_{N0}|$, (c) total energy error $|\mathcal{E}_N - \mathcal{E}_{N0}$, phase space discretization N, time step Δt, regularization 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 (–): $(N, \Delta t, \epsilon) = (512^2, 0.25, 0.1)$.

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.

4.2. Two-stream instability

There are several commonly studied forms of the two-stream instability. Following [29, 31], we consider the initial distribution

$$f_0(x, v) = v^2 f_M(v; v_{\text{th}}) (1 + \alpha \cos(kx)), \quad (x, v) \in [0, 4\pi] \times [-6, 6], \tag{32}$$

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 20 shows the electron phase space density. The first row shows times 0,3,6, and 9 highlighting the formation of the central vortex characteristic of

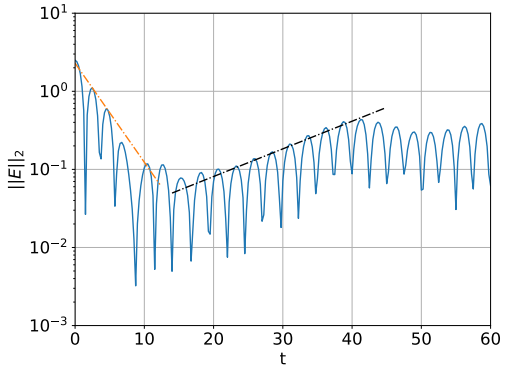


Fig. 15: Strong Landau damping, 2-norm of electric field vs. time, phase space discretization $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, run time 754 s, decay rate ($\gamma = -0.2920$, orange line), growth rate ($\gamma = 0.0815$, black line) [33].

DISTRIBUTION A: Distribution approved for public release.

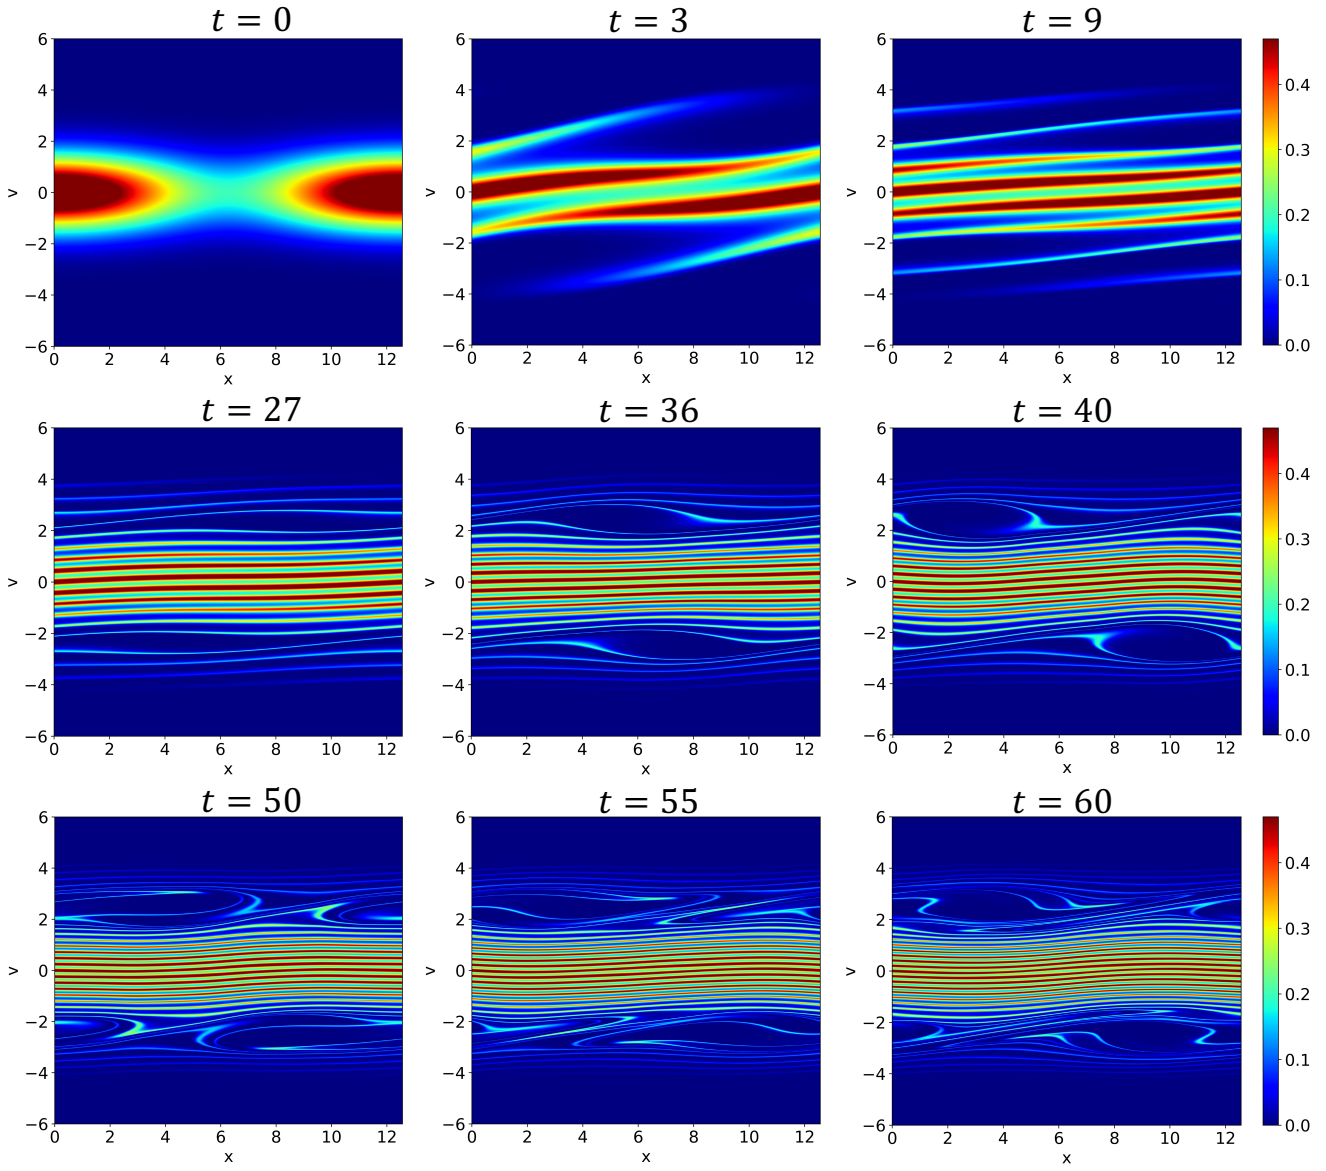


Fig. 16: Strong Landau damping, electron distribution $f(x, v, t)$, discretization $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, run time 754 s.

the two-stream instability. The second row shows times 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 near velocity $v = \pm 3$, caused by the entrainment of low density phase fluid, and the propagation of these vortices from one period to the next.

The run time of the AMR computation was 437 s and the corresponding uniform calculation took 745 s. The key difference compared to the Landau damping simulations is the formation of a large central vortex. However above and below the vortex, the two streams have some features similar to those seen in strong Landau damping; the entrainment of low density phase fluid to form small vortices near velocity $v = \pm 3$, and propagation of the small vortices 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. 21. Fig. 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

DISTRIBUTION A: Distribution approved for public release.

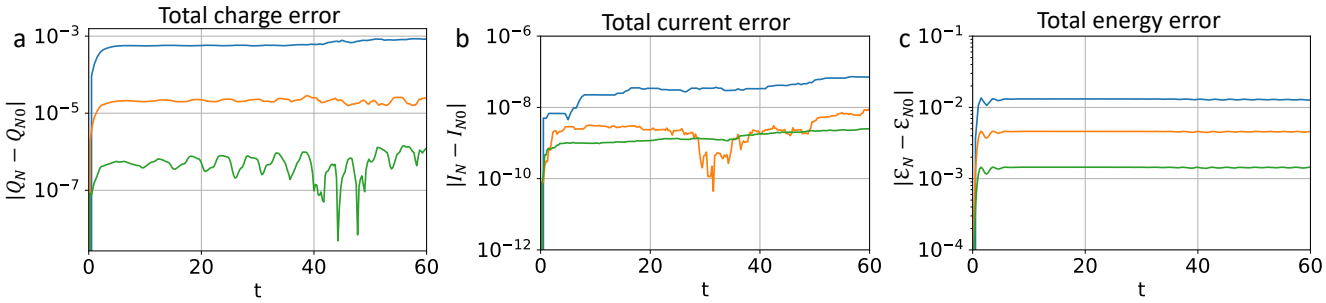


Fig. 17: Strong Landau damping, conservation properties, BLTC parameters MAC $\theta = 0.6$, degree $n = 8$, mesh size N, regularization parameter ϵ, time step Δt refined simultaneously, (a) total charge error $|Q_N - Q_{N0}|$, (b) total current error $|J_N - J_{N0}|$, (c) total energy error $|\mathcal{E}_N - \mathcal{E}_{N'}|$, 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

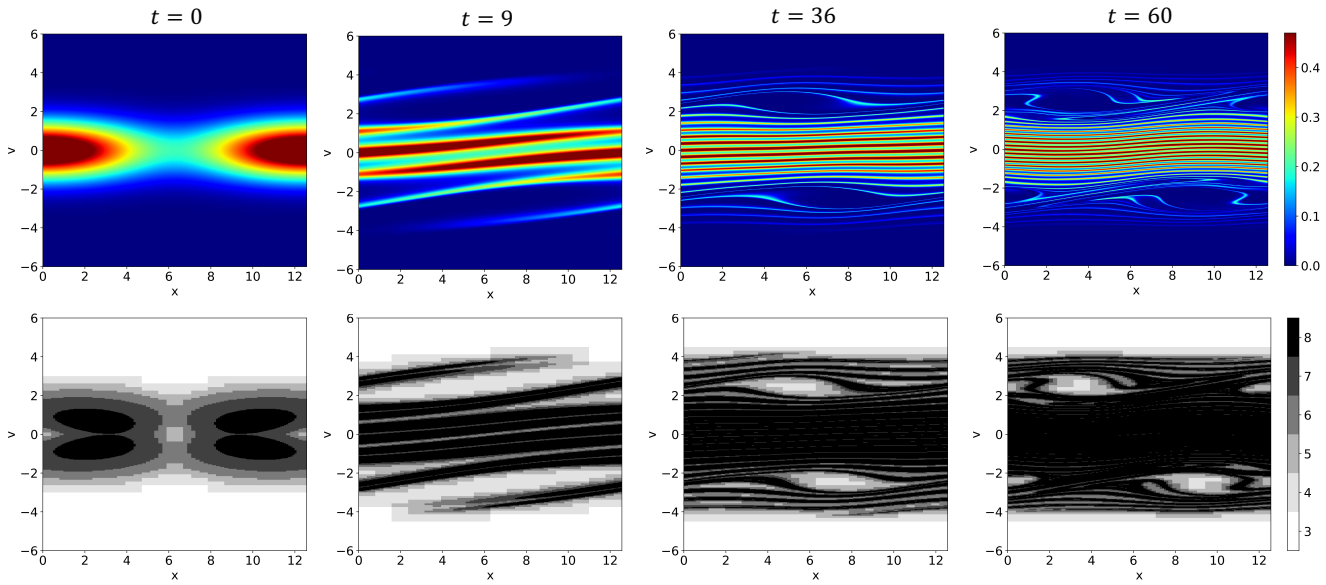


Fig. 18: Strong Landau damping, AMR simulation, $t_f = 60$, $N = 16 \times 64 + 5$ AMR levels, $N_0 = 154297$, $N_f = 538601$, or 14.7% to 51.2% of the number of points in the uniform calculation, comparable uniform resolution is $N = 512 \times 2048$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, $\epsilon_{AMR} = 0.01$, GPU time 345 s, left: phase space, right: panel level.

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.

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,

$$f_0(x, v) = \frac{1}{2}\left(f_M(v - v_s; v_{th}) + f_M(v + v_s; v_{th})\right)(1 + \alpha \cos kx), \quad (x, v) \in [0, 4\pi] \times [-6, 6]. \tag{33}$$

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. 22. Times $t = 0$ to $t = 10$ show acceleration of some particles to velocities 2.5× 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 $t = 14$

DISTRIBUTION A: Distribution approved for public release.

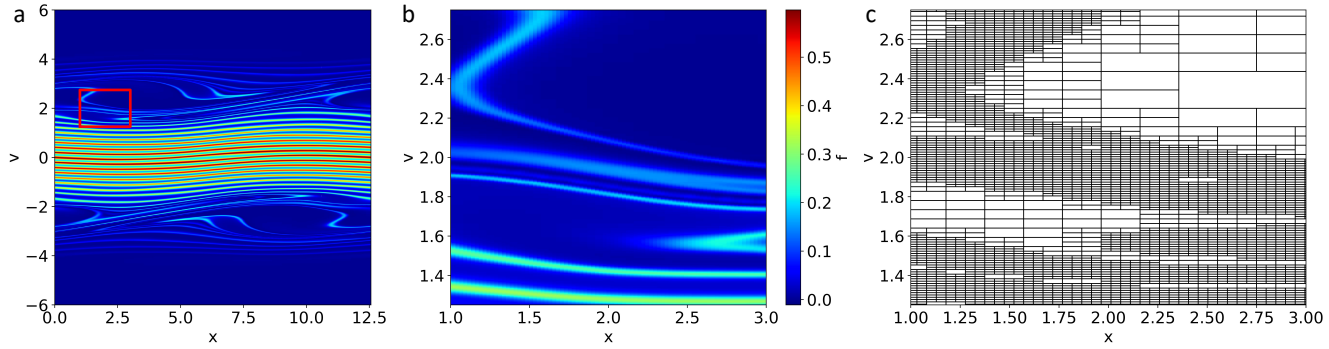


Fig. 19: Strong Landau damping, AMR simulation, $N = 16\times64 + 5$ AMR levels, comparable uniform resolution is $N = 512\times2048$, regularization parameter $\epsilon = 0.1$, time step $\Delta t = 0.25$, AMR tolerance $\epsilon_{AMR} = 0.01$, run time 345 s, (a) phase space at time $t = 60$, (b) zoom of small region in phase space centered at $(x, v) = (2, 2)$, (c) phase space panels in zoom region of (b).

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. 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. 23. Fig. 23(a) shows the uniform phase-space with phase-space discretization $N = 512\times1024$, (b) shows the AMR phase-space with discretization $N = 16\times32 + 5$ 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.

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

$$\partial_s f + x'\partial_x f + \left(\frac{q}{m\gamma_b^3\beta_b^2c^2}E - \kappa x\right)\partial_{x'} f = 0, \quad E = -\partial_x\phi, \quad -\partial_x^2\phi = \frac{q}{\epsilon_0}n, \tag{34}$$

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 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,

$$-\partial_x\phi(s, x) \to \pm\frac{q}{2\epsilon_0}N_i, \quad x \to \mp\infty, \tag{35}$$

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

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

$$f^{eq}(x, x') = \frac{\hat{n}}{\sqrt{2\pi T^*}}\exp\left(-\frac{x'^2}{2T^*}\right)\tilde{n}_\Delta\left(\frac{x}{\gamma_b\lambda_D}\right), \tag{36}$$

where $\hat{n}$ is the peak density, T^* is the dimensionless temperature, λ_D is the Debye length, $\hat{\omega}_p$ is the plasma frequency, Δ is a dimensionless parameter,

$$\hat{n} = \frac{\epsilon_0 m\gamma_b^3\beta_b^2c^2\kappa}{q^2(1+\Delta)}, \quad T^* = \frac{T}{m\gamma_b\beta_b^2c^2}, \quad \lambda_D = \left(\frac{T}{m\hat{\omega}_p^2}\right)^{1/2}, \quad \hat{\omega}_p = \left(\frac{q^2\hat{n}}{\epsilon_0 m}\right)^{1/2}, \quad \Delta = \frac{\gamma_b^3\beta_b^2c^2\kappa}{\hat{\omega}_p^2} - 1, \tag{37}$$

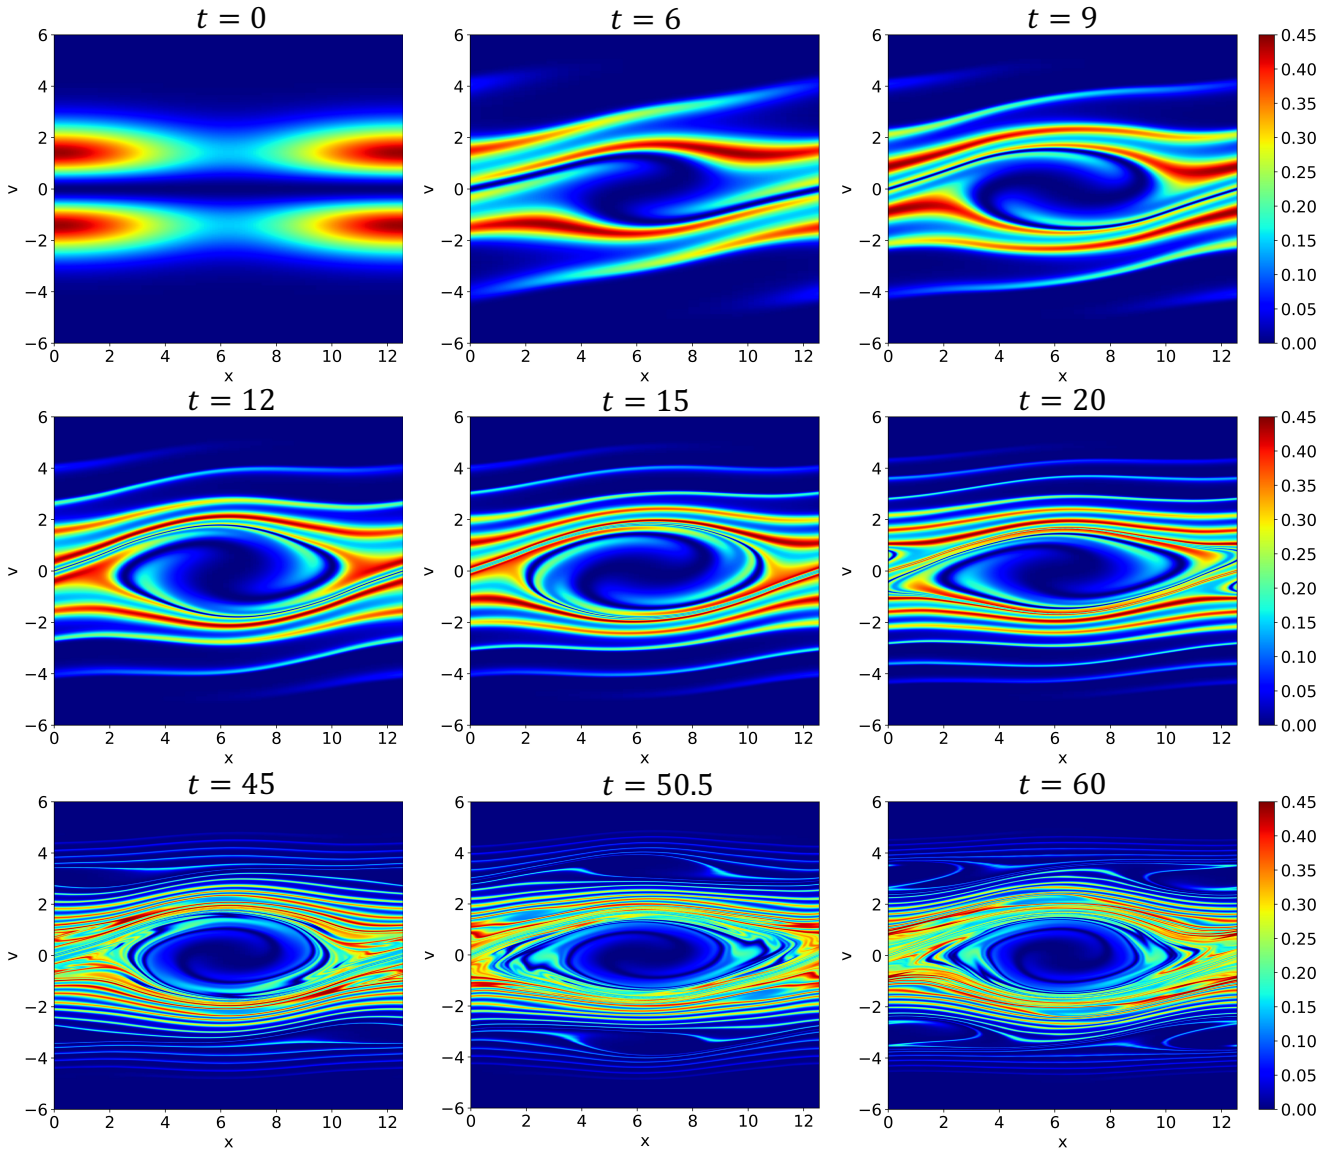


Fig. 20: Two-stream instability, phase space electron distribution, initial condition Eq. (32), discretization $N = 16{\times}64 + 5$ AMR levels, $N = 264325$ initially, which is 25.1% of the uniform N and $N = 639165$ ultimately which is 60.8% of the uniform N, regularization parameter $\epsilon = 0.1$, run time 437 s compared to uniform run time 739 s.

and the scaled density $\tilde{n}_\Delta$ satisfies the following equation and normalizing conditions,

$$\tilde{n}''_\Delta + (1+\Delta)\tilde{n}_\Delta - \tilde{n}_\Delta^2 + \frac{1}{\tilde{n}_\Delta}\left(\tilde{n}'_\Delta\right)^2 = 0, \quad \tilde{n}(0) = 1, \quad \tilde{n}'_\Delta(0) = 0. \tag{38}$$

We solved Eq. (38) numerically using the parameter values in Table 2 [4], and Fig. 24 shows the resulting equilibrium distribution which has a smooth top-hat profile. The connection of the scaled density $\tilde{n}_\Delta$ 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,

$$\psi''_\Delta(\hat{x}) = 1 + \Delta - e^{-\psi_\Delta(\hat{x})}\,, \tag{39}$$

satisfying $\psi'_\Delta(0) = \psi_\Delta(0) = 0$. Identifying $\tilde{n}_\Delta = e^{-\psi_\Delta}$, differentiating, and substituting Eqn. (39) leads to the Eqn. (38) for $\tilde{n}_\Delta$.

At this point it is convenient to scale variables,

$$x \to \kappa^{1/2}x, \quad \phi \to \frac{q}{m\gamma_b^3\beta_b^2c^2}\phi, \quad f \to \frac{q^2\kappa}{m\gamma_b^3\beta_b^2c^2\epsilon_0}f, \tag{40}$$

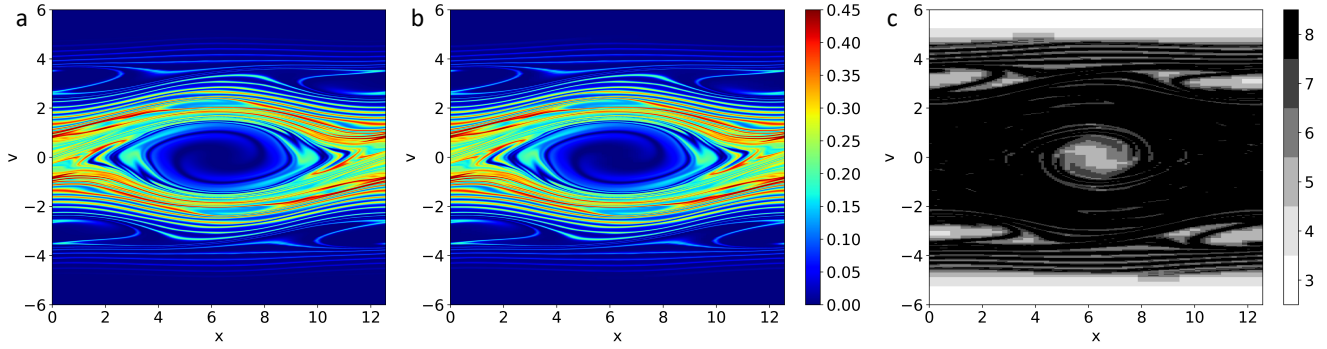


Fig. 21: Two-stream instability, time $t = 60$, time step $\Delta t = 0.25$, regularization parameter $\epsilon = 0.1$, compare AMR with uniform resolution, (a) uniform discretization, $N = 528 \times 2048$, run time 739 s, (b) AMR, $\epsilon_{AMR} = 0.008$, $N = 16 \times 64 + 5$ AMR levels, equivalent to 25.1% to 60.8% of uniform N, run time 435 s, (c) AMR panels shaded by level.

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

after which the Vlasov-Poisson equations are

$$\partial_s f + x' \partial_x f + (E - x)\, \partial_{x'} f = 0, \quad E = -\partial_x \phi, \quad -\partial_x^2 \phi = n, \quad n = \int_{-\infty}^{\infty} f dx' \,, \tag{41}$$

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

$$f(x, x', s = 0) = f^{eq}(x/\mu, \mu x'), \quad \mu = 1.25, \tag{42}$$

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

$$E(s, x) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} K_\epsilon(x, y) f(y, x', s) dy dx', \quad K_\epsilon(x, y) = \frac{1}{2} \frac{x - y}{((x - y)^2 + \epsilon^2)^{1/2}}, \tag{43}$$

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

The calculation used phase-space domain $|x| \le \sqrt{\kappa}\ 15$ mm $= \pi/100$ and $|x'| \le 14.5$ mrad, discretization $N = 16^2 + 6$ AMR levels, axial increment $\Delta s = \sqrt{\kappa}\, L_p/10 = \pi/3$, regularization parameter $\epsilon = \sqrt{\kappa}\, 0.1$ mm $\approx 2.09 \times 10^{-4}$, and BLTC pararameters MAC $\theta = 0.8$ and degree $n = 4$. Figure 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 26 presents the solution at axial position $s = \sqrt{\kappa}\, 40 L_p = 40\pi/3$ with (a) uniform discretization $N_u = 1024^4$, (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 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× faster, using on average 5× fewer points. The AMR calculation requires only 15% of the uniform field evaluation time and one third of the uniform remeshing time.

DISTRIBUTION A: Distribution approved for public release.

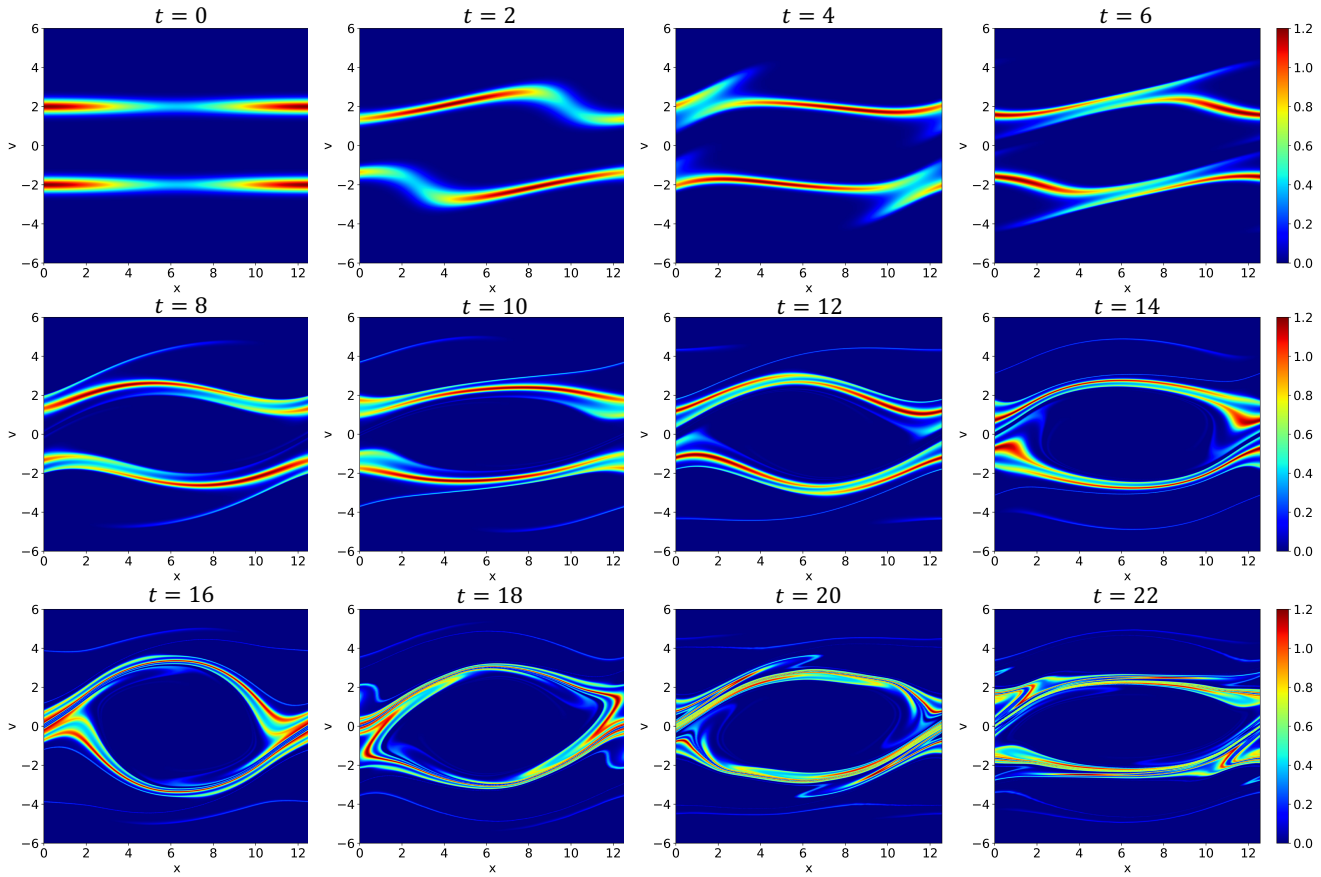


Fig. 22: Cold two-stream instability, initial distribution Eq. (33), phase space discretization $N = 16 \times 32 + 5$ AMR levels, $N_0 = 129245$ which is 24.6% of the uniform N, N_f is 193345, which is 36.8% of he uniform N, AMR tolerance $\epsilon_{AMR} = 0.008$, regularization parameter $\epsilon = 0.1$, time step $\Delta t = 0.25$, run time time 44.1 s.

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
16^2+ 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 3: Halo formation in a mismatched beam, profile of results shown in Fig. 26 discretization N, initial and final values N_0, N_f, average value N_{ave}, total run time, field evaluation time, remeshing time

5. Summary

A forward semi-Lagrangian scheme called FARSIGHT was presented 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 [44].

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 [13], and the cloudy

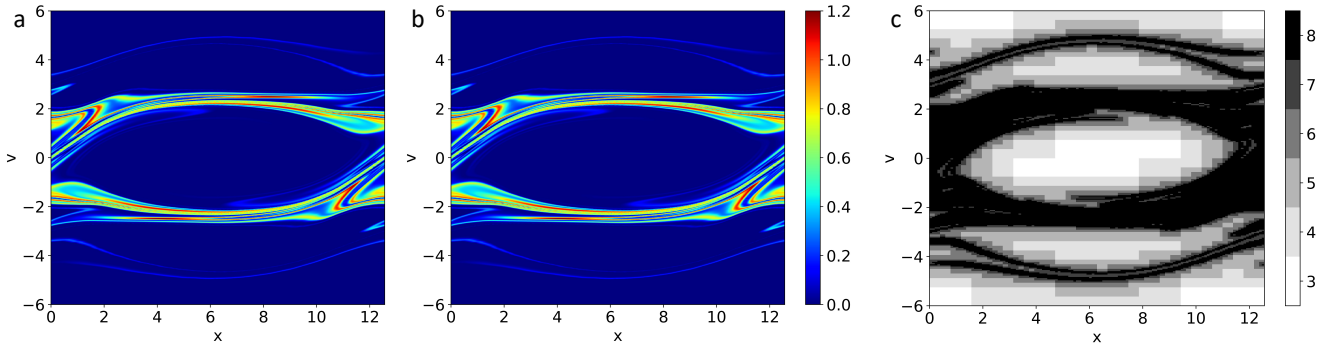


Fig. 23: Cold two-stream instability, compare AMR with uniform resolution at time in 22, $\epsilon = 0.1, \Delta t = 0.25$ (a) $N = 528 \times 1024$, GPU time 123.5 s, (b) $N = 16 \times 32 + 5$ levels of AMR, $\epsilon_{AMR} = 0.008$, GPU time 44.1 s, (c) AMR resolution.

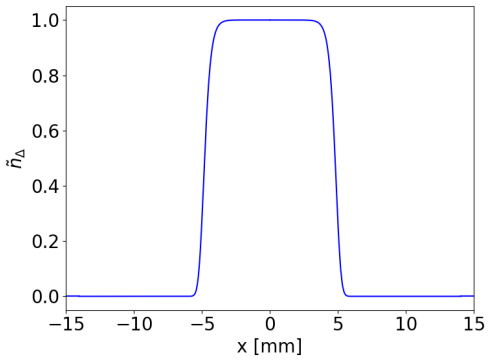


Fig. 24: Scaled density $\tilde{n}_\Delta$ for equilibrium sheet beam distribution [3, 4], $\Delta = 5.522 \times 10^{-8}$.

Vlasov solver [37] 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 [63].

Acknowledgments

This work was supported by AFOSR grant FA9550-19-1-0072. Computational resources were provided by Advanced Research Computing at the University of Michigan.

References

[1] B. Svedung Wettervik, T. C. Dubois, T. Fülöp, Vlasov modelling of laser-driven collisionless shock acceleration of protons, Phys. Plasmas 23 (2016).
[2] A. Grassi, L. Fedeli, A. Sgattoni, A. Macchi, Vlasov simulation of laser-driven shock acceleration and ion turbulence, Plasma Phys. Contr. F. 58 (2016).
[3] S. M. Lund, A. Friedman, G. Bazouin, Sheet beam model for intense space charge: Application to Debye screening and the distribution of particle oscillation frequencies in a thermal equilibrium beam, Phys. Rev. Spec. Top.-Ac 14 (2011) 054201.
[4] M. Campos Pinto, E. Sonnendrücker, A. Friedman, D. P. Grote, S. M. Lund, Noiseless Vlasov-Poisson simulations with linearly transformed particles, J. Comput. Phys. 275 (2014) 236–256.
[5] E. Sonnendrucker, J. J. Barnard, A. Friedman, D. P. Grote, S. M. Lund, Simulation of heavy ion beams with a semi-Lagrangian Vlasov solver, Nucl. Instrum. Meth. A 464 (2001) 470–476.

DISTRIBUTION A: Distribution approved for public release.

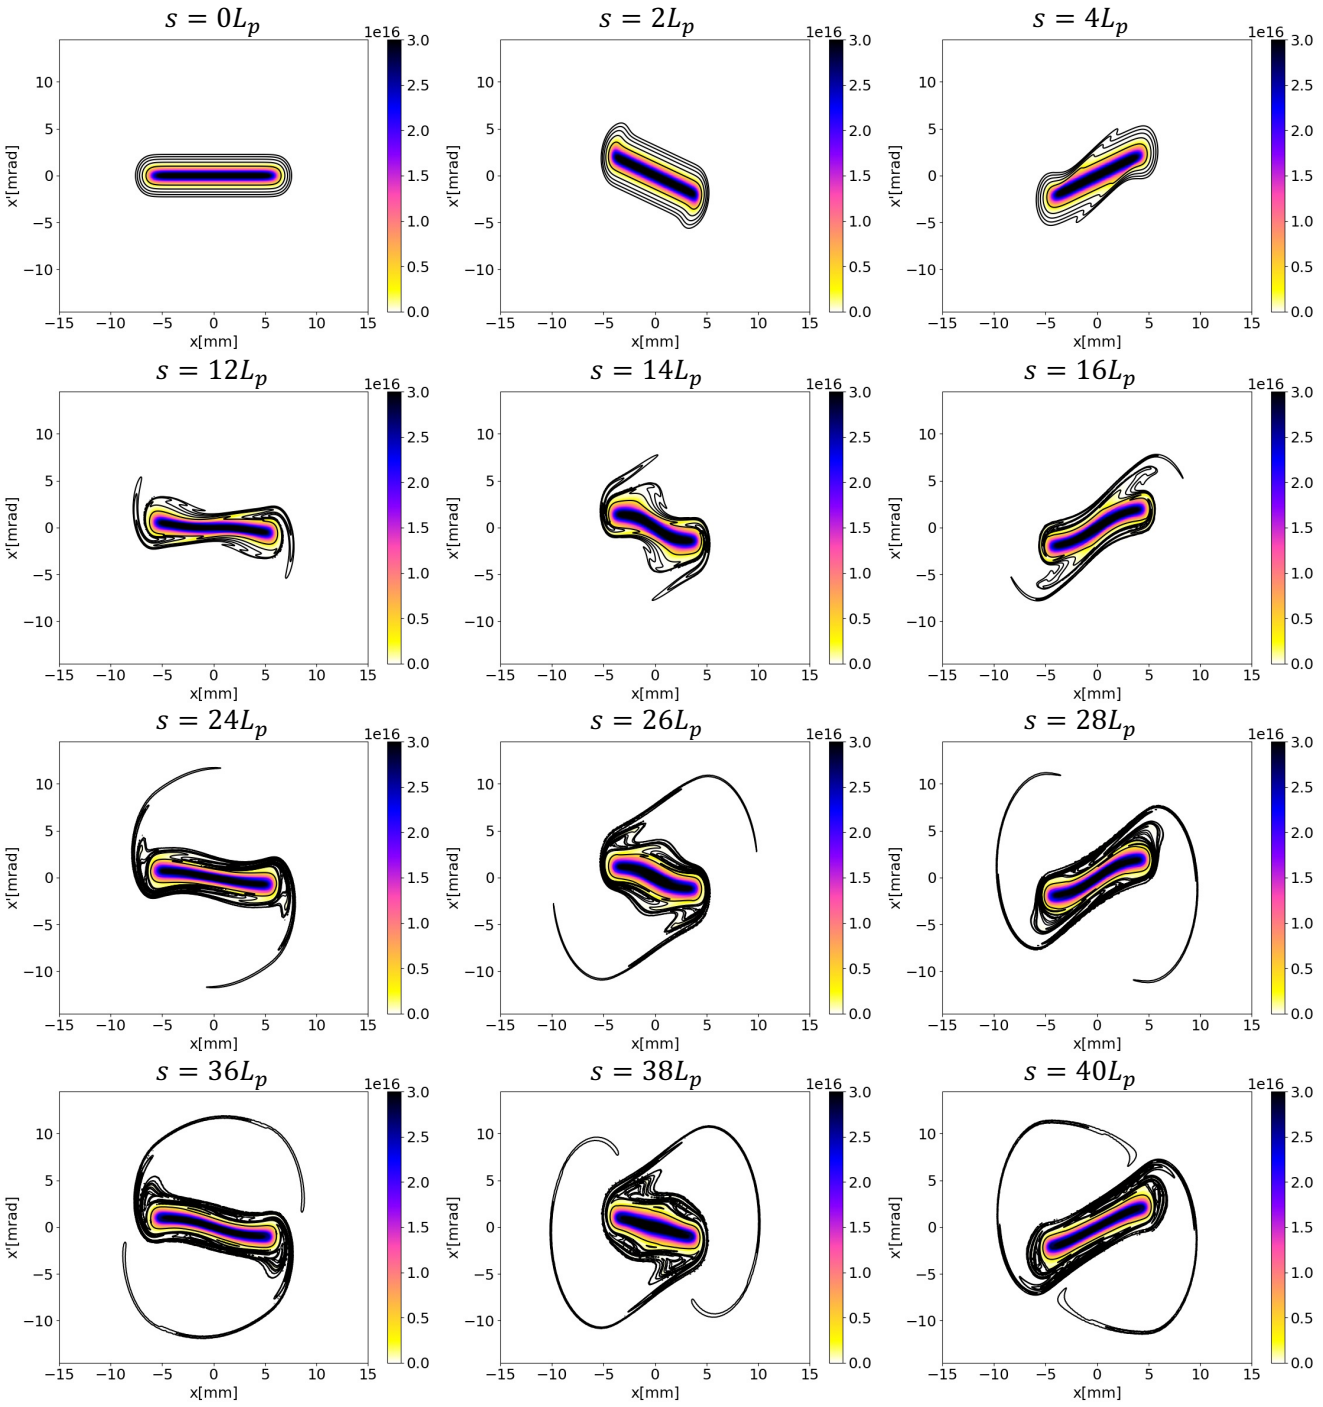


Fig. 25: Mismatched beam problem, phase space discretization $N = 16^2 + 6$ AMR levels, axial increment $\Delta s = 0.1L_p$, regularization parameter $\epsilon = 0.1$ mm, GPU run time 245 s, results agree well with reference solution in [4].

[6] B. Afeyan, F. Casas, N. Crouseilles, A. Dodhy, E. Faou, M. Mehrenberger, E. Sonnendrücker, Simulations of kinetic electrostatic electron nonlinear (KEEN) waves with variable velocity resolution grids and high-order time-splitting, Eur. Phys. J. D 68 (2014) 295.
[7] B. A. Shadwick, G. M. Tarkenton, E. Esarey, C. B. Schroeder, Fluid and Vlasov models of low-temperature, collisionless, relativistic plasma interactions, Phys. Plasmas 12 (2005) 056710.
[8] A. G. Thomas, Vlasov simulations of thermal plasma waves with relativistic phase velocity in a Lorentz boosted frame, Phys. Rev. E 94 (2016).
[9] C. K. Birdsall, A. B. Langdon, Plasma Physics via Computer Simulation, IOP Publishing, 1991.
[10] R. Hockney, J. Eastwood, Computer Simulation Using Particles, Taylor & Francis, 1988.
[11] W. M. Nevins, G. W. Hammett, A. M. Dimits, W. Dorland, D. E. Shumaker, Discrete particle noise in particle-in-cell simulations of plasma

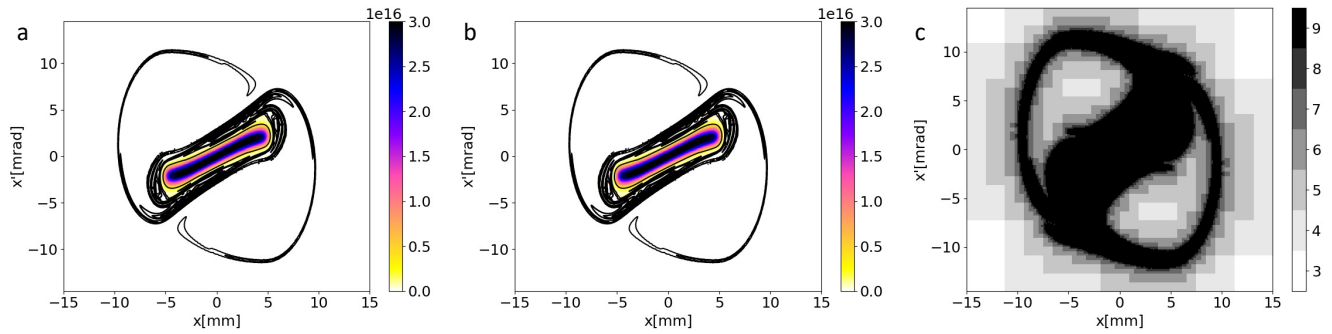


Fig. 26: Mismatched beam problem, phase space at axial position $s = 40L_p$, (a) uniform discretization, $N = 1024^2$, run time 1169 s, (b) AMR, $N = 16^2 + 6$ AMR levels, initially $N = 135253$ or 12.9% of N_u, finally $N = 207721$ or 19.8% of N_u, run time 245 s, (c) panels in AMR calculation.

microturbulence, Phys. Plasmas 12 (2005) 1–16.
[12] D. W. Hewett, Fragmentation, merging, and internal dynamics for PIC simulation with finite size particles, J. Comput. Phys. 189 (2003) 390–426.
[13] D. J. Larson, C. V. Young, A finite mass based method for Vlasov-Poisson simulations, J. Comput. Phys. 284 (2015) 171–185.
[14] A. Christlieb, R. Krasny, J. Verboncoeur, J. Emhoff, I. Boyd, Grid-free plasma simulation techniques, IEEE T. Plasma Sci. 34 (2006) 149–165.
[15] D. M. Thomas, J. T. Holgate, A treecode to simulate dust–plasma interactions, Plasma Physics and Controlled Fusion 59 (2016) 025002.
[16] T. Kessler, S. Rjasanow, S. Weisser, Vlasov-Poisson system tackled by particle simulation utilizing boundary element methods, SIAM J. Sci. Comput. 42 (2020) B299–B326.
[17] P. I. Shustov, I. V. Kuzichev, I. Y. Vasko, A. V. Artemyev, A. J. Gerrard, The dynamics of electron holes in current sheets, Phys. Plasmas 28 (2021).
[18] M. V. Goldman, D. L. Newman, P. Pritchett, Vlasov simulations of electron holes driven by particle distributions from PIC reconnection simulations with a guide field, Geophys. Res. Lett. 35 (2008).
[19] L. Li, J. Qiu, G. Russo, A high-order semi-lagrangian finite difference method for nonlinear vlasov and bgk models, Comm. App. Math. Com. Sc. (2022) 1–29.
[20] J. W. Banks, J. A. F. Hittinger, A new class of nonlinear finite-volume methods for Vlasov simulation, IEEE T. Plasma Sci. 38 (2010) 2198–2207.
[21] G. Vogman, U. Shumlak, P. Colella, Conservative fourth-order finite-volume Vlasov–Poisson solver for axisymmetric plasmas in cylindrical (r,v,v) phase space coordinates, J. Comput. Phys. 373 (2018).
[22] R. Heath, I. Gamba, P. Morrison, C. Michler, A discontinuous Galerkin method for the Vlasov–Poisson system, J. Comput. Phys. 231 (2012).
[23] J. T. Parker, P. J. Dellar, Fourier–Hermite spectral representation for the Vlasov–Poisson system in the weakly collisional limit, J. Plasma Phys. 81 (2015).
[24] E. Camporeale, G. Delzanno, B. Bergen, J. Moulton, On the velocity space discretization for the Vlasov–Poisson system: Comparison between implicit Hermite spectral and Particle-in-Cell methods, Comput. Phys. Commun. 198 (2016).
[25] F. Filbet, T. Xiong, Conservative discontinuous galerkin/hermite spectral method for the vlasov–poisson system, Communications on Applied Mathematics and Computation (2020) 1–26.
[26] J. Denavit, Numerical simulation of plasmas with periodic smoothing in phase space, J. Comput. Phys. 9 (1972) 75–98.
[27] C. Z. Cheng, G. Knorr, The Integration of the Vlasov Equation in Configuration Space, J. Comput. Phys. 22 (1976) 330–351.
[28] E. Sonnendrücker, J. Roche, P. Bertrand, A. Ghizzo, The Semi-Lagrangian Method for the Numerical Resolution of the Vlasov Equation, J. Comput. Phys. 149 (1999) 201–220.
[29] J. A. Rossmanith, D. C. Seal, A positivity-preserving high-order semi-Lagrangian discontinuous Galerkin scheme for the Vlasov-Poisson equations, J. Comput. Phys. 230 (2011) 6203–6232.
[30] J. M. Qiu, C. W. Shu, Positivity preserving semi-Lagrangian discontinuous Galerkin formulation: Theoretical analysis and application to the Vlasov-Poisson system, J. Comput. Phys. 230 (2011) 8386–8409.
[31] Y. Cheng, I. M. Gamba, P. J. Morrison, Study of conservation and recurrence of Runge-Kutta discontinuous Galerkin schemes for Vlasov-Poisson systems, J. Sci. Comput. 56 (2013) 319–349.
[32] C. S. Huang, T. Arbogast, C. H. Hung, A semi-Lagrangian finite difference WENO scheme for scalar nonlinear conservation laws, J. Comput. Phys. 322 (2016) 559–585.
[33] D. Sirajuddin, W. N. Hitchon, A truly forward semi-Lagrangian WENO scheme for the Vlasov-Poisson system, J. Comput. Phys. (2019).
[34] A. Christlieb, M. Link, H. Yang, R. Chang, High-Order Semi-Lagrangian WENO Schemes Based on Non-polynomial Space for the Vlasov Equation, Commun. Appl. Math. Comput. (2021).
[35] X. Cai, W. Guo, J. M. Qiu, A high order semi-Lagrangian discontinuous Galerkin method for Vlasov–Poisson simulations without operator splitting, J. Comput. Phys. 354 (2018) 529–551.
[36] T. Xiong, G. Russo, J. M. Qiu, Conservative Multi-dimensional Semi-Lagrangian Finite Difference Scheme: Stability and Applications to the Kinetic and Fluid Simulations, J. Sci. Comput. 79 (2019) 1241–1270.
[37] C. Alard, S. Colombi, A cloudy Vlasov solution, Mon. Not. R. Astron. Soc. 359 (2005) 123–163.
[38] S. Colombi, C. Alard, A 'metric' semi-Lagrangian Vlasov–Poisson solver, J. Plasma Physics. 83 (2017).
[39] B. Wang, G. H. Miller, P. Colella, A particle-in-cell method with adaptive phase-space remapping for kinetic plasmas, SIAM J. Sci. Comput. 33 (2011) 3509–3537.
[40] J. Hittinger, J. Banks, Block-structured adaptive mesh refinement algorithms for Vlasov simulation, J. Comput. Phys. 241 (2013).

DISTRIBUTION A: Distribution approved for public release.

[41] M. Gutnic, M. Haefele, I. Paun, E. Sonnendrücker, Vlasov simulations on an adaptive phase-space grid, Comput. Phys. Commun. 164 (2004) 214–219.
[42] N. Besse, E. Deriaz, Madaule, Adaptive multiresolution semi-Lagrangian discontinuous Galerkin methods for the Vlasov equations, J. Comput. Phys. 332 (2017) 376–417.
[43] W. Guo, Y. Cheng, An adaptive multiresolution Discontinous Galerkin method for time-dependent transport equations in multidimensions, SIAM J. Sci. Comput. 39 (2017) A2962–A2992.
[44] N. Vaughn, L. Wilson, R. Krasny, A GPU-Accelerated Barycentric Lagrange Treecode, in: 21st IEEE Int. Workshop Parallel Distrib. Sci. Eng. Comput. (PDSEC), 2020, pp. 701–710.
[45] P. Bosler, L. Wang, Christiane Jablonowski, R. Krasny, A Lagrangian particle/panel method for the barotropic vorticity equations on a rotating sphere, Fluid Dyn. Res. 46 (2014).
[46] A. J. Chorin, Numerical study of slightly viscous flow, J. Fluid Mech. 57 (1973) 785–796.
[47] R. Krasny, Desingularization of periodic vortex sheet roll-up, J. Comput. Phys. 65 (1986) 292–313.
[48] G. Cottet, P. Koumoutsakos, Vortex Methods: Theory and Practice, Cambridge University Press, 2000.
[49] M. Perlman, On the accuracy of vortex methods, J. Comput. Phys. 59 (1985) 200–223.
[50] P. Koumoutsakos, Inviscid axisymmetrization of an elliptical vortex, J. Comput. Phys. 138 (1997) 821–857.
[51] A. Myers, P. Colella, B. Van Straalen, A 4th-order particle-in-cell method with phase-space remapping for the Vlasov-Poisson equation, SIAM J. Sci. Comput. 39 (2017) B467–B485.
[52] J. Barnes, P. Hut, A hierarchical O(N log N) force-calculation algorithm, Nature 324 (1986) 446–449.
[53] L. Wang, R. Krasny, S. Tlupova, A kernel-independent treecode based on barycentric Lagrange interpolation, Commun. Comput. Phys. 28 (2020) 1415–1436.
[54] J.-P. Berrut, L. N. Trefethen, Barycentric Lagrange interpolation, SIAM Rev. 46 (2004) 501–517.
[55] L. Landau, On the vibrations of the electron plasma, volume 10, 1946, pp. 25–34.
[56] C. Villani, Particle systems and nonlinear Landau damping, Phys. Plasmas 21 (2014) 030901.
[57] J. T. Beale, A. Majda, Vortex Methods. II: Higher Order Accuracy in Two and Three Dimensions, Math. Comput. 39 (1982).
[58] P. A. Raviart, An analysis of particle methods, in: F. Brezzi (Ed.), Numerical Methods in Fluid Dynamics, Springer Berlin Heidelberg, Berlin, 1985, pp. 243–324.
[59] N. Crouseilles, M. Mehrenberger, E. Sonnendrücker, Conservative semi-Lagrangian schemes for Vlasov equations, J. Comput. Phys. 229 (2010) 1927–1953.
[60] J. M. Qiu, C. W. Shu, Conservative semi-Lagrangian finite difference WENO formulations with applications to the Vlasov equation, Commun. Comput. Phys. (2011).
[61] X. Cai, J. Qiu, J. M. Qiu, A conservative semi-Lagrangian HWENO method for the Vlasov equation, J. Comput. Phys. 323 (2016) 95–114.
[62] X. Cai, S. Boscarino, J.-M. Qiu, High order semi-lagrangian discontinuous galerkin method coupled with runge-kutta exponential integrators for nonlinear vlasov dynamics, J. Comput. Phys. 427 (2021) 110036.
[63] L. Wilson, N. Vaughn, R. Krasny, A GPU-accelerated fast multipole method based on barycentric Lagrange interpolation and dual tree traversal, Comput. Phys. Commun. 265 (2021) 108017.

DISTRIBUTION A: Distribution approved for public release.

Project 3: Use of plasma waves to generate intense XUV radiation

The following is an earlier preprint of the paper that was published in Phys. Rev. Lett.

DISTRIBUTION A: Distribution approved for public release.

Unlimited Photon Acceleration

R. Sandberg* and A. G. R. Thomas†
Gérard Mourou Center for Ultrafast Optical Sciences, University of Michigan, Ann Arbor, MI 48109, USA
(Dated: January 5, 2022)

Plasma wakefield accelerators have been subject of substantial recent interest because of their extremely high acceleration gradients for electrons. The propagating density gradients of the wakefield may also be used to frequency upshift a trailing witness laser pulse, a process known as 'photon acceleration'. In a uniform plasma, the witness laser will eventually dephase from the accelerating gradient because of group delay in the plasma. Here, we find the phase matching conditions for the pulse using a tapered density profile. An analytic solution for a 1d nonlinear plasma wake with an electron beam driver indicates that, even though the plasma density decreases, the frequency shift reaches no asymptotic limit, i.e. is unlimited provided the wake can be sustained. In quasi-3d particle-in-cell simulations using the calculate density profile we demonstrate frequency shifts of the witness laser beam up to ten times, limited only by simulation resolution and non-optimal driver evolution. In 1d simulations, more than 30 times shifts were observed. The pulse energy increases in this process, by a factor of five, the wake guides the pulse and it is temporally compressed by group velocity dispersion, resulting in the resulting XUV laser pulse having relativistic ($a_0 \sim 1$) intensity.

The many applications of bright, coherent extreme ultraviolet (XUV) light have motivated substantial interest in source development, such as the construction of XUV wavelength Free Electron Lasers such as FLASH [1] as well as nonlinear frequency mixing [2], high harmonic generation [3], and XUV lasing [4], to name a few. Another method for generating short wavelengths is 'photon acceleration' [5, 6]. Dressed by a background plasma, 'quasi-photons' gain an effective mass $\hbar\omega_p/c^2$, where $\omega_p^2 = e^2 n/m_e\varepsilon_0$ 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 [7].

Photon acceleration can arise as a result of plasma wakefields [5], ionization fronts [8, 9] and even using metamaterials [10]. Photon acceleration was measured in ionization front [11] and laser wakefield acceleration experiments [12, 13]. Recent results include cascaded sequences of localized ionizations [14], resulting in large frequency shifts and the use of plasma wakes to downshift radiation to very long wavelengths [15]. Limits to photon acceleration in plasma wakefields in the linear [16] and nonlinear regimes [17] were previously studied, identifying dephasing of the photon beam with respect to the accelerating refractive index gradient 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 [18]. The flying focus was also used to mitigate the analogous process of electron beam dephasing in a plasma wakefield [19], in addition to related spatiotemporally structured focusing schemes [20, 21].

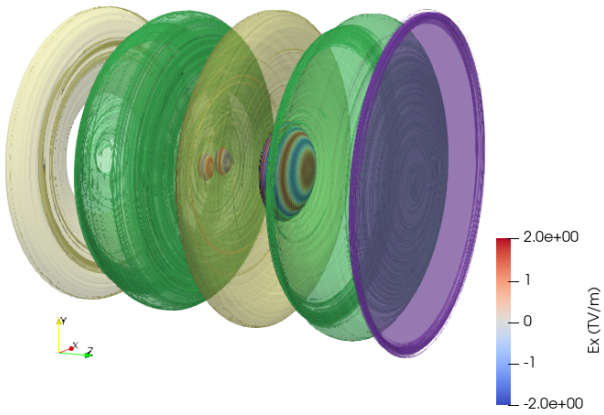


FIG. 1: Isocontours from quasi-3d simulation data demonstrating phase matched photon acceleration; the purple disk at the right is the drive electron beam, yellow and green are the positive and negative density perturbations of the wake, and the blue and red spheroid is a contour of the envelope of the transverse electric field, colored to show the laser pulse field phase.

Another method for mitigating dephasing, in the context of electron acceleration is the use of tapered plasma density ramps [22–25]. 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

DISTRIBUTION A: Distribution approved for public release.

shifts in photon acceleration [5, 16].

In this Letter, we develop 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, Fig. 1. The pulse experiences net energy gain of 5× to 250 mJ and is compressed, leading to an ultrashort pulse of XUV radiation with relativistic intensity.

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 (photon kinetic) solutions can be used for the temporal variation in the light [6, 26, 27] to relate it to the density gradients generated in a wakefield. For a given dispersion relation D, the frequency of an optical mode ω propagating in the z direction with wavenumber k_z will vary with distance propagated as

$$\frac{d\omega}{dz} = \frac{\partial D/\partial t}{\partial D/\partial k_z} \,. \tag{1}$$

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 $(x, y, z, \zeta = t - \int_0^z dz'/v_d(z'))$ [23]. The change in the frequency can, therefore, be expressed as

$$\frac{d\omega}{dz} = \frac{\partial D/\partial \zeta}{\partial D/\partial k_z} \,. \tag{2}$$

For example, for linear plasma dispersion, $D = \omega^2 - \omega_p^2/\gamma - k^2c^2$, 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 and can be ignored,

$$\frac{d\omega}{dz} \simeq \frac{1}{2k_z c^2}\frac{\partial(\omega_p^2/\gamma)}{\partial \zeta} \,. \tag{3}$$

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.

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 a tapered density profile (similar to that proposed for mitigating dephasing in electron acceleration [23, 25, 28, 29]) to continuously increase the plasma wavelength and keep the laser pulse experiencing a positive plasma density gradient.

For convenience, we choose the point where the density perturbation in the wake $\delta n = 0$ within the positive gradient region, 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 on 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 $\zeta_{\rm centroid} = \zeta_\delta$ for all times, where $\zeta_{\rm centroid}$ denotes the center of the witness laser pulse. Expressed in differential form, we have

$$\frac{d\zeta_{\rm centroid}}{dz} = \frac{d\zeta_\delta}{dz} \,, \tag{4}$$

which may be written as

$$\frac{1}{v_{\rm centroid}} - \frac{1}{v_d} = \frac{d\zeta_d}{dn}\frac{dn}{dz} \,. \tag{5}$$

The quantity n refers to the unperturbed plasma density ahead of the driver and is a function of z. For an ultrarelativistic particle beam driver, we make the approximation $v_d \to c$. 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,

$$\frac{dn}{dz} \simeq \frac{1}{2c}\frac{\omega_p^2}{\omega^2}\left[\frac{d\zeta_\delta}{dn}\right]^{-1} \,. \tag{6}$$

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

Here, we derive the phase-matching condition for a one-dimensional (1d) nonlinear plasma wake, which we find is also accurate to describe the three-dimensional case for a broad driver ($k_p\sigma_r \gg 1$), where $k_p = \omega_p/c$ and σ_r is the is the root-mean-square beam radius. 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, we write the perturbed density of the wake $n_w(\zeta; z)$ as a function of ζ parameterized by z, i.e. by the local density $n(z)$. Solutions for a plasma

DISTRIBUTION A: Distribution approved for public release.

wake generated by a square temporal-profile beam-driver in a uniform density plasma may be expressed parametrically [30, 31] such that the amplitude and the starting phase of the wake behind the driver are functions of the drive beam density n_d, length L_d and the unperturbed plasma density n. The amplitude is described by the maximum Lorentz factor $\gamma_m = 1 - n_d\phi_d$, where ϕ_d is the electrostatic potential at the end of the drive beam. The starting phase of the wake behind the driver is defined by

$$\tan\varphi_{st} = \frac{-E_d}{p_d}\sqrt{\frac{\gamma_d - 1}{2}}\,, \tag{7}$$

where E_d, p_d, and γ_d are the electric field, momentum, and Lorentz factor of the plasma at the end of the drive beam. These quantities are obtained from solutions to the fluid equations with the drive beam density included [31]. The position ζ_δ, measured from the *leading* edge of the drive beam, may be expressed as

$$\zeta_\delta = \zeta(3\pi/2) - \zeta(\varphi_{st}) + L_d/c\,, \tag{8}$$

where $\zeta(\varphi)$ is the position in the undriven wake as a function of a periodic parametric coordinate φ, with $\varphi(0)$ corresponding to a density maximum. The function $\zeta(\varphi)$ is defined implicitly by

$$\omega_p\zeta(\varphi) = \frac{2}{\kappa'}E(\varphi,\kappa) - \kappa' F(\varphi,\kappa) - \frac{2\kappa}{\kappa'}\sin\varphi\,, \tag{9}$$

where $F(\phi,\kappa)$ and $E(\varphi,\kappa)$ are elliptic integrals of the first and second kind and

$$\kappa^2 = \frac{\gamma_m - 1}{\gamma_m + 1}, \quad \kappa'^2 = 1 - \kappa^2\,. \tag{10}$$

Hence, the change in the location of the phase-matched refractive index gradient used in Eqn. 6, $d\zeta_\delta/dn$, is defined implicitly using Eqns. 7 and 8. The plasma wake density gradient at ζ_δ is given by

$$\left.\frac{\partial n_w}{\partial\zeta}\right|_{\zeta=\zeta_\delta} = \omega_p n\sqrt{2\left[\gamma_m - 1\right]}\,. \tag{11}$$

For clarity, we express the equations in normalized units, with $n \to n/n_0$ and $n_w(\zeta;z) \to n_w(\zeta;z)/n_0$, where we use n_0 as a reference plasma density, $\omega_p \to \omega_p/\omega_{p0}$, and $\omega \to \omega/\omega_{p0}$ etc., where ${\omega_{p0}}^2 = e^2 n_0/m_e\varepsilon_0$. In normalized units, combining Eqns. 3 and 11 yields

$$\frac{d\omega^2}{dz} \simeq n^{3/2}\sqrt{2\left[\gamma_m - 1\right]}\,, \tag{12}$$

and Eqn. 6 is

$$\frac{dn}{dz} \simeq \frac{1}{2}\frac{n}{\omega^2}\left[\frac{d\zeta_\delta}{dn}\right]^{-1}\,. \tag{13}$$

Eqns. (12) and (13) 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 of arbitrary length and density. Finding the variation of ζ_δ with density in the nonlinear case is a challenge as it is the solution of an implicit equation. Insight can be gained, however, by using an ultrashort driver approximation, for which an analytic solution can be found.

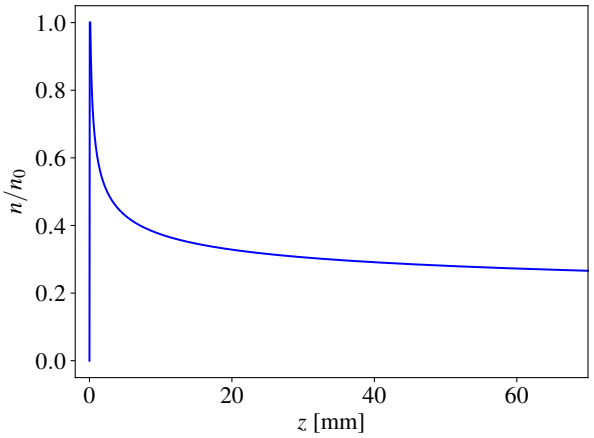


FIG. 2: Density profile for unlimited photon acceleration driven by an ultrashort driver with $A = (k_{p0}L_{d0})(n_{d0}/n_0) = 0.38$.

The ultrashort driver approximation means fixing $A = n_d L_d$ and letting $L_d \to 0$, so that $n_d = A/L_d \to \infty$. In this limit, the maximum amplitude of the wave reduces to $\gamma_m(n;A) = 1 + A^2/2n$ and the phase is $\varphi_{st} = \pi/2$.

Eqns. (12) and (13) may be combined and directly integrated to obtain

$$\omega(n) = \omega_0\exp\left[A\left(\zeta_\delta(n) - \zeta_{\delta 0}\right)\right] \tag{14}$$

defined for n in $(0, n_0]$ 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-Bendixon theorem, there are no limit sets to the orbits. In fact, as $n \to 0$, $\zeta_\delta(n) \to \infty$ and $\omega(n) \to \infty$. Hence, we call this model 'unlimited photon acceleration.' We can determine the evolution of the system in terms of propagation distance z via the implicitly defined equation for $n(z)$,

$$z = 2\omega_0^2\int_{n_0}^{n}\frac{1}{n'}\frac{d\zeta_\delta}{dn'}e^{2A(\zeta_\delta(n) - \zeta_{\delta 0})}dn'\,. \tag{15}$$

For a moderate strength driver, $A < 1$, the location of the zero density perturbation, ζ_δ, can be described accurately by the expansion

$$\zeta_\delta = \frac{\pi}{n^{1/2}} + \frac{2A}{n} + \frac{3\pi}{16}\frac{A^2}{n^{3/2}} + \cdots\,. \tag{16}$$

The density profile for $A = 0.4$ is plotted in figure 2. The shape is characteristic for all values of A, with a sharp

initial decrease in density that rolls over, but always decreases monotonically.

We use the quasi-3d particle-in-cell code FBPIC [32] to demonstrate the phase-matched photon acceleration in practice. 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 [33]. 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 $c\tau = 1.25\ \mu$m, and a spot size of $w_0 = 2\lambda_{p0} = 20\lambda_{L0} = 16\ \mu$m, starting with its centroid $\zeta_{\rm cent} = 6.6\ \mu$m/c behind the front of the electron drive beam. For these parameters, $\zeta_\delta = 5.4\ \mu$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 [28]. Visualization of the drive beam, its wake and the witness laser pulse in the quasi-3d simulation is shown in Fig. 1.

We use the plasma density profile shown in Fig. 2 and determined by equations (14) and (15), scaled by $n_0 = 1.74 \times 10^{19}$ cm^{-3}, which is 100× 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 PMPA can be seen in Fig. 3. (a) shows the spectrum as a function of length propagated in the simulation. Overlaid is the theoretical model of Eqns (14-16) (red line), which predicts the shift of the photons well. (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. The drive beam loses 25 J of energy over the course of the simulation 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 [34], and so the energy gain by the pulse would be expected to scale with the frequency increase. In the simulations, the nearly 10× frequency shift results in a lower 5× 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

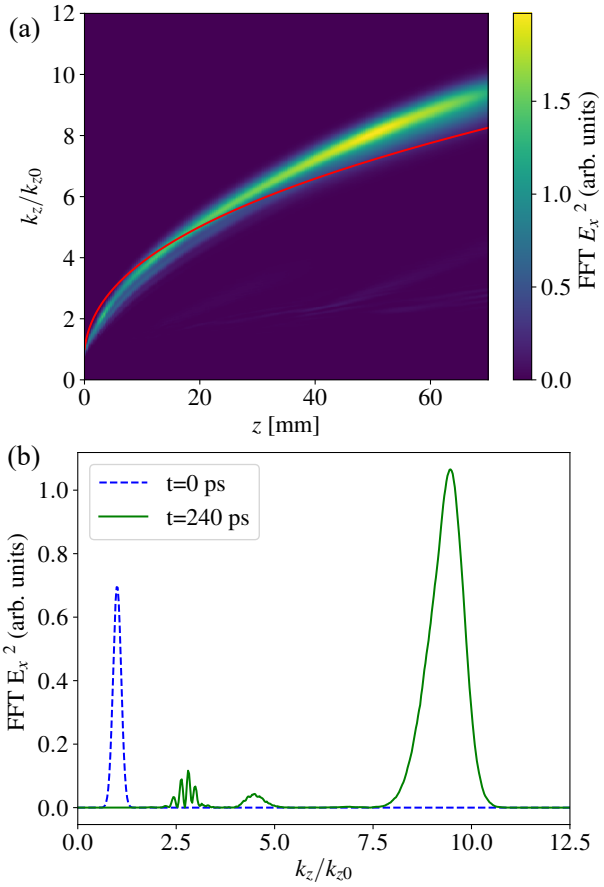


FIG. 3: Quasi-3d simulation for a drive beam with $k_{p0}L_{d0} = 0.5$. (a) Laser spectrum as a function of time. Shown in red is the theoretical curve given by Eqns (14-16) (b) Initial and final spectra.

pulse develops significant up-chirp, which, together with dispersion, leads to a 2.7× compression of the pulse from an initial duration of $c\tau = 1.25\ \mu$m to 0.51 μm. The intensity sees a significant increase of 20× from an initial intensity $I = 2 \times 10^{18}$ W/cm^2 to 4×10^{19} 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 $z_R = \pi w_0^2/\lambda_{Lf} \approx 10$ mm, meaning that the pulse 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

DISTRIBUTION A: Distribution approved for public release.

longer drive pulse, where tuning of the drive beam propagation was easier through parametric studies of many simulations, we have demonstrated even 30× frequency shifts. These 1d results will be published in a longer follow up publication. 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 from light, which may be useful for photolithography.

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 arbitary frequency shift limited only by the drive beam. In the quasi-3d simulations shown here, an unoptimized single short 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 [35].

Quasi-3d simulation results directly demonstrated a 10× frequency shift with 1% energy transfer efficiency. The wake itself serves as a wave guide. Combined with relativistic self-focusing and the lengthening of the Rayleigh range as the frequency increases, the pulse is guided for the duration of the interaction. In quasi-3d simulations not shown here (to be published in a longer follow-up publication), 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× 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, for example using a 'flying focus' [18] to mitigate issues with drive laser dispersion.

This work was supported by AFOSR grant FA9550-19-1-0072 and NSF grant 1804463. Computational resources were provided by Advanced Research Computing at the University of Michigan.

* ryansand@umich.edu
† agrt@umich.edu

[1] 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. ISSN 0370-1573. doi: https://doi.org/10.1016/j.physrep.2019.02.002. URL https://www.sciencedirect.com/science/article/pii/S0370157319300663. 10 years of pioneering X-ray science at the Free-Electron Laser FLASH at DESY.

[2] 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. doi: 10.1088/1367-2630/abae89. URL https://doi.org/10.1088/1367-2630/abae89.

[3] 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. doi: 10.1103/PhysRevLett.79.2967. URL https://link.aps.org/doi/10.1103/PhysRevLett.79.2967.

[4] 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. doi:10.1103/PhysRevLett.81.5804. URL https://link.aps.org/doi/10.1103/PhysRevLett.81.5804.

[5] S. C. Wilks, J. M. Dawson, W. B. Mori, T. Katsouleas, and M. E. Jones. Photon accelerator. *Physical Review Letters*, 62(22), 5 1989. ISSN 0031-9007.

[6] Jose Tito Mendonca. *Theory of photon acceleration*. Institute of Physics Pub, 2001. ISBN 0750307110.

[7] 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. ISSN 10797114. doi: 10.1103/PhysRevLett.96.014803.

[8] E. Esarey, G. Joyce, and P. Sprangle. Frequency upshifting of laser pulses by copropagating ionization fronts. *Physical Review A*, 44(6):3908–3911, 9 1991. ISSN 10502947. doi:10.1103/PhysRevA.44.3908.

[9] 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. ISSN 00933813. doi:10.1109/27.509995.

[10] 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. doi: 10.1038/s41467-019-09313-8. URL https://doi.org/10.1038/s41467-019-09313-8.

[11] 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. doi:10.1103/PhysRevLett.78.4773. URL https://link.aps.org/doi/10.1103/PhysRevLett.78.4773.

[12] 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. doi: 10.1103/PhysRevLett.76.3570. URL https://link.aps.org/doi/10.1103/PhysRevLett.76.3570.

[13] 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. Mendonca, 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. doi:

DISTRIBUTION A: Distribution approved for public release.