

UCLA

UCLA Electronic Theses and Dissertations

Title

On the Acceleration and Transport of Electrons Generated by Intense Laser-Plasma Interactions at Sharp Interfaces

Permalink

<https://escholarship.org/uc/item/6cb3h11p>

Author

May, Joshua Joseph

Publication Date

2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

On the Acceleration and Transport
of Electrons Generated by Intense Laser-Plasma Interactions
at Sharp Interfaces

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

by

Joshua Joseph May

2017

© Copyright by
Joshua Joseph May
2017

ABSTRACT OF THE DISSERTATION

On the Acceleration and Transport of Electrons Generated by Intense Laser-Plasma Interactions at Sharp Interfaces

by

Joshua Joseph May

Doctor of Philosophy in Physics

University of California, Los Angeles, 2017

Professor Warren B. Mori, Chair

The continued development of the chirped pulse amplification technique has allowed for the development of lasers with powers of in excess of $10^{15}W$, for pulse lengths with durations of between .01 and 10 picoseconds, and which can be focused to energy densities greater than 100 giga-atmospheres. When such lasers are focused onto material targets, the possibility of creating particle beams with energy fluxes of comparable parameters arises. Such interactions have a number of theorized applications. For instance, in the Fast Ignition concept for Inertial Confinement Fusion [1], a high-intensity laser efficiently transfers its energy into an electron beam with an appropriate spectra which is then transported into a compressed target and initiate a fusion reaction. Another possible use is the so called Radiation Pressure Acceleration mechanism, in which a high-intensity, circularly polarized laser is used to create a mono-energetic ion beam which could then be used for medical imaging and treatment, among other applications. For this latter application, it is important that the laser energy is transferred to the ions and not to the electrons. However the physics of such high energy-density laser-matter interactions is highly kinetic and non-linear, and presently not fully understood.

In this dissertation, we use the Particle-in-Cell code OSIRIS [2, 3] to explore the generation and transport of relativistic particle beams created by high intensity lasers focused onto solid density matter at normal incidence. To explore the generation of relativistic electrons by such interactions, we use primarily one-dimensional (1D) and two-dimensional (2D), and a few three-dimensional simulations (3D). We initially examine the idealized case of normal incidence of relatively short, plane-wave lasers on flat, sharp interfaces. We find that in 1D the results are highly dependent on the initial temperature of the plasma, with significant absorption into relativistic electrons only possible when the temperature is high in the direction parallel to the electric field of the laser. In multi-dimensions, absorption into relativistic electrons arises independent of the initial temperature for both fixed and mobile ions, although the absorption is higher for mobile ions. In most cases however, absorption remains at 10's of percent, and as such a standing wave structure from the incoming and reflected wave is setup in front of the plasma surface. The peak momentum of the accelerated electrons is found to be $2a_0m_ec$, where $a_0 \equiv eA_0/m_ec^2$ is the normalized vector potential of the laser in vacuum, e is the electron charge, m_e is the electron mass, and c is the speed of light. We consider cases for which $a_0 > 1$. We therefore call this the $2a_0$ acceleration process. Using particle tracking, we identify the detailed physics behind the $2a_0$ process and find it is related to the standing wave structure of the fields. We observe that the particles which gain energy do so by interacting with the laser electric field within a quarter wavelength of the surface where it is at an anti-node (it is a node at the surface). We find that only particles with high initial momentum – in particular high transverse momentum – are able to navigate through the laser magnetic field as its magnitude decreases in time each half laser cycle (it is an anti-node at the surface) to penetrate a quarter wavelength into the vacuum where the laser electric field is large. For a circularly polarized laser the magnetic field amplitude never decreases at the surface, instead its direction simply rotates. This prevents electrons from leaving the plasma and they therefore cannot gain energy from the electric field.

For pulses with longer durations ($\gtrsim 250fs$), or for plasmas which do not have initially sharp interfaces, we discover that in addition to the $2a_0$ acceleration at the surface, relativistic particles are also generated in an underdense region in front of the target. These particles

have energies without a sharp upper bound. Although accelerating these particles removes energy from the incoming laser, and although the surface of the plasma does not stay perfectly flat and so the standing wave structure becomes modified, we find in most cases, the $2a_0$ acceleration mechanism occurs similarly at the surface and that it still dominates the overall absorption of the laser.

To explore the generation of relativistic electrons at a solid surface and transport of the heat flux of these electrons in cold or warm dense matter, we compare OSIRIS simulations with results from an experiment performed on the OMEGA laser system at the University of Rochester. In that experiment, a thin layer of gold placed on a slab of plastic is illuminated by an intense laser. A greater than order-of-magnitude decrease in the fluence of hot electrons is observed when those electrons are transported through a plasma created from a shock-heated plastic foam, as compared to transport through cold matter (unshocked plastic foam) at somewhat higher density. Our simulations indicate two reasons for the experimental result, both related to the magnetic field. The primary effect is the generation of a collimating B-field around the electron beam in the cold plastic foam, caused by the resistivity of the plastic. We use a Monte Carlo collision algorithm implemented in OSIRIS to model the experiment. The incoming relativistic electrons generate a return current. This generates a resistive electric field which then generates a magnetic field from Faraday's law. This magnetic field collimates the forward moving relativistic electrons. The collisionality of both the plastic and the gold are likely to be greater in the experiment than the 2D simulations where we used a lower density for the gold (to make the simulations possible) which heats up more. In addition, the use of 2D simulations also causes the plastic to heat up more than expected. We compensated for this by increasing the collisionality of the plasma in the simulations and this led to better agreement. The second effect is the growth of a strong, reflecting B-field at the edge of the plastic region in the shock heated material, created by the convective transport of this field back towards the beam source due to the neutralizing return current. Both effects appear to be caused primarily by the difference in density in the two cases. Owing to its higher heat capacity, the higher density material does not heat up as much from the heat flux coming from the gold, which leads to a larger resistivity.

Lastly, we explored a numerical effect which has particular relevance to these simulations, due to their high energy and plasma densities. This effect is caused by the use of macro particles (which represent many real particles) which have the correct charge to mass ratio but higher charge. Therefore, any physics of a single charge that scales as q^2/m will be artificially high. Physics that involves scales smaller than the macro-particle size can be mitigated through the use of finite size particles. However, for relativistic particles the spatial scale that matters is the skin depth and the cell sizes and particle sizes are both smaller than this. This allows the wakes created by these particles to be artificially high which causes them to slow down much faster than a single electron. We studied this macro-particle stopping power theoretically and in OSIRIS simulations. We also proposed a solution in which particles are split in to smaller particles as they gain energy. We call this effect Macro Particle Stopping. Although this effect can be mitigated by using more particles, this is not always computationally efficient. We show how it can also be mitigated by using high-order particle shapes, and/or by using a particle-splitting method which reduces the charge of only the most energetic electrons.

The dissertation of Joshua Joseph May is approved.

Troy A. Carter

Chandra J. Joshi

Christoph Niemann

Warren B. Mori, Committee Chair

University of California, Los Angeles

2017

*To Zoë,
whose patience, understanding, and love,
supported me on this long journey*

TABLE OF CONTENTS

1	Introduction	1
1.1	Fast Ignition Inertial Confinement Fusion	3
1.2	Accepted and Proposed Mechanisms for Light Absorption	11
1.3	Particle-in-Cell Codes and OSIRIS	18
1.4	Collisions in PIC	22
1.5	Resistive Collimation	37
1.6	Outline of Dissertation	40
2	High Intensity Laser Absorption in One- and Two-Dimensions	43
2.1	1D Simulations	44
2.1.1	Fixed Ions	45
2.1.2	Mobile Ions	56
2.2	Early Time 2D Simulations - Absorption at a Flat Surface	61
2.2.1	Preamble - Early Time Evolution of and Absorption by an Initially Cold Plasma	62
2.2.2	Particle Tracking in Semi-1D - Acceleration Mechanism at Sharp Surfaces	65
2.2.3	Single Particle Model and Numerical Calculations	80
2.2.4	Initial Heating Mechanism - ‘Slipping Stream’ Instability	90
2.2.5	Circular Polarization, and s-Polarization	97
2.3	Absorption in the Hydrodynamic Regime - Picosecond Timescales and Micron-Scale Preplasmas	101
2.3.1	Hydrodynamic Evolution of an Initially Flat Surface	102
2.3.2	Plasma Evolution and Laser Absorption in the Presence of Preplasma	121

2.3.3	Circular and Elliptical Polarization	135
2.4	Conclusion	140
3	Modeling of Electron Transport Experiment on OmegaEP	143
3.1	Experimental Setup and Results	145
3.1.1	Target design	145
3.1.2	Plasma Characterization	146
3.1.3	Experimental Results of Electron Transport	148
3.2	Simulations	150
3.2.1	Challenges and Limitations	150
3.2.2	Simulation Setup	154
3.2.3	Simulation Results	158
3.3	Summary	175
4	Macro-Particle Stopping in Particle-in-Cell Simulations	177
4.1	Introduction	177
4.2	Theory	179
4.3	Single Particle Stopping in Fast Ignition Simulations	189
4.4	Relaxation of High-Energy Tails	193
4.5	Integrated Fast Ignition Simulations	195
4.6	Particle Splitting	196
4.7	Summary	202
5	Conclusion	204
A	Conservation of Transverse Canonical Momentum	210

B	High-Energy Particle Damper in OSIRIS	212
B.1	Refluxing	212
B.2	Damper	215
B.3	Fourth-root Temperature Calculation	219
C	Analytic Model for $2a_0m_e c$ Peak Momentum	221
	References	224

LIST OF FIGURES

1.1	PIC loop	20
2.1	Phase space for 1D PIC with $T_{e_{initial}} = 0$	47
2.2	Phase space for 1D PIC with $T_{e_{initial}} = 5.0keV$	49
2.3	Phase space for 1D PIC with $T_{e_{initial}} = 316keV$	51
2.4	Phase space for 1D PIC with $T_{e_{initial}} = 5keV$ only in x_1	53
2.5	Laser absorption vs. u_{th}	55
2.6	p_1x_1 for 1D cold with mobile ions	57
2.7	Heat flux in e^- and ions	59
2.8	p_1x_1 ion phase space for 1D plasma	60
2.9	p_1x_1 phase space for 2D plasma at early times	64
2.10	Heat flux at 55 fs	66
2.11	Phase spaces in semi-1D	68
2.12	Tracks of accelerated particles in position space	70
2.13	Phase space tracks and single particle model	76
2.14	Tracks of $v \cdot E$	79
2.15	Single particle skin depth effects	83
2.16	Single particle laser absorption effects	85
2.17	Finite skin depth using the model of GE Kemp	88
2.18	Bootstrap heating	94
2.19	Surface rippling	96
2.20	s-Polarized light phase space	98
2.21	Circularly polarized light and RPA	100
2.22	Electron spectrum at 300fs	105

2.23	Density at 300fs	106
2.24	Selection rules for tracking high γ electrons	110
2.25	Tracks at $t\omega_0 = 1540$	111
2.26	Tracks at $t\omega_0 = 1555$	112
2.27	Tracks at $t\omega_0 = 1570$	113
2.28	Tracks at $t\omega_0 = 1585$	114
2.29	Tracks at $t\omega_0 = 1600$	115
2.30	Expansion into plasma over 3 picoseconds	117
2.31	2D Surface evolution	119
2.32	Heat flux spectrum	121
2.33	ρ, q_1, j_1 vs. $\log_{10}(\gamma - 1)$	122
2.34	Expansion of initial preplasma	124
2.35	Density and heat flux with preplasma	125
2.36	Heat flux spectrum vs. $\log_{10}(\gamma - 1)$	126
2.37	Heat flux spectrums for $10\mu m$ density scale lengths	130
2.38	Magnetic field formed with $10\mu m$ pre-plasma	133
2.39	Surfaces under linearly, elliptically, and circularly polarized laser illumination	137
2.40	Density and heat flux with circular, elliptical polarization	139
3.1	Omega EP experimental setup	146
3.2	Al absorption spectrum	147
3.3	Rad-hydro simulation results	149
3.4	Results from eXport-10B experiment	151
3.5	eXport10B simulation initial density	154
3.6	Heat flux for $7n_c$ and $63n_c$	159

3.7	Forward heat flux spectrum in gold	160
3.8	B_3 for $7n_c$ and $63n_c$	161
3.9	Pressure and transverse heat flux	163
3.10	q_1 for 7 and 63 n_c for 0.026 and 40 eV	165
3.11	q_1 and B_3 for varied collisionality in $63n_c$ plastic	168
3.12	q_1 and B_3 for varied collisionality in $28n_c$ plastic	169
3.13	q_1 and B_3 for varied collisionality in $7n_c$ plastic	171
4.1	Theoretical wakes for different size and shape particles	184
4.2	Reduction in stopping force due to finite-sized particles	186
4.3	Wake in 2D PIC simulation	191
4.4	Relaxation of high-energy tails	194
4.5	Heating and power in FI target core	197
4.6	Particle count and heatflux in FI-relevant simulation	199
4.7	Particle spectrum and heatflux without and with particle splitting	201
5.1	Momentum phase spaces in 3D simulations	207
B.1	Boundary condition effect on phase space	216

ACKNOWLEDGMENTS

First and foremost I would like to thank my advisor, Warren Mori, for all the knowledge, support, and guidance he has given me over the years. It goes without saying that no one person has given me more help in completing this dissertation, and all the research that went into it.

I would further like to thank my collaborators from UC San Diego, Farhat Beg, Ming-Sheng Wei, and Christopher McGuffey, for helpful discussions with regards to the material in Chapter 3. And I would also like to thank my collaborators from (at the time) IST in Lisbon, Ricardo Fonseca, Luís O. Silva, and most especially Frederico Fiúza, for many discussions on a broad array of topics.

Also, I would like to thank my fellow students and the other members of the Plasma Simulation group at UC Los Angeles, Adam Tableman, Peicheng Yu, Weiming An, Michail Tzoufras, Benjamin Winjum, Frank Tsung, Viktor Decyk, and Asher Davidson. It has been a pleasure working and learning with you all.

Lastly I would like to thank my parents, Karen and Dennis, who have supported me without condition in whatever I have chosen to pursue.

VITA

1999-2004	BS with Highest Honors, Physics, BA with Honors, Mathematics, UC Santa Cruz, Santa Cruz, CA
2000-2003	Medicinal Chemistry Intern, Berlex Biosciences, Richmond, CA – Setup, executed, and diagnosed organic reactions in pursuit of novel pharmacological (anti-coagulant) compounds.
2004-Present	Ph.D. Candidate, Physics and Astronomy, UC Los Angeles, Westwood, Los Angeles, CA
2004-2008	Teaching Assistant, UC Los Angeles, Westwood, Los Angeles, CA – Assisted in the teaching of both lower-division lab and upper division lecture classes.
2008-Present	Graduate Student Researcher, UC Los Angeles, Westwood, Los Angeles, CA – Assisted in the development of the Particle-in-Cell code OSIRIS; and used that code to conduct research on high-mass-density, high-energy-density, laser-matter interactions.
2013-Present	Treasurer and Board Member, Urban Soil-Tierra Urbana, a 501(c)3 Limited-Equity Housing Coop, Koreatown, Los Angeles, CA

PUBLICATIONS

J. May, J. Tonge, I. Ellis, W. B. Mori, F. Fiúza, R. A. Fonseca, L. O. Silva, and C. Ren
Enhanced stopping of macro-particles in particle-in-cell simulations, Physics of Plasmas Vol.
21 (2014), p.052703

J. May, J. Tonge, F. Fiúza, R. A. Fonseca, L. O. Silva, C. Ren, and W. B. Mori, *Mechanism of generating fast electrons by an intense laser at a steep overdense interface*, Physical Review E, Vol. 84 No.2 (2011), p.025401

Frederico Fiúza, Ricardo A. Fonseca, Luís O. Silva, John Tonge, Joshua May, and Warren B. Mori, *Three-Dimensional Simulations of Laser-Plasma Interactions at Ultrahigh Intensities*, IEEE Transactions on Plasma Science, Vol. 39, No. 11 (2011), p.2618

F. Fiúza, M. Marti, R. A. Fonseca, L. O. Silva, J. Tonge, J. May, and W. B. Mori, *Efficient modeling of laser-plasma interactions in high energy density scenarios*, Plasma Phys. Control. Fusion Vol. 53 (2011) p. 074004

J. Tonge, J. May, W. B. Mori, F. Fiúza, S. F. Martins, R. A. Fonseca, L. O. Silva, and C. Ren, *A simulation study of fast ignition with ultrahigh intensity lasers*, Physics of Plasmas Vol. 16 (2009), p.056311

CHAPTER 1

Introduction

In one form or another, the technology of human society will eventually rely primarily if not solely on fusion energy — the only question is whether that fusion happens in our sun or on our planet. Indeed it could even be said that we are already relying on fusion energy, as almost all¹ energy sources on Earth have their source in the fusion reactions inside the sun, whether it be the heat from sunlight today which drives the evaporation and rain which powers hydroelectric plants; or the sunlight from billions of years ago which grew the organisms which eventually became oil and coal.

However the point is that there is a key difference between those two power sources just mentioned – although the rain will (hopefully) keep falling for the indefinite future – were we to keep using fossil fuels, eventually they would be depleted. Indeed it is even more stark than that, as the international consensus on carbon fossil fuels being the cause of global warming means humanity will almost certainly have to stop using them with plenty left in the ground.

Therefore, as we inevitably move away from fossil fuels, we have two (not necessarily conflicting) options. One, we can increase the scale and efficiency with which we harvest energy from solar radiation, possibly through hydroelectric² but more likely through photovoltaics, solar thermal power plants, wind turbines, biofuels, et cetera. Or, two, we can ‘bring the Sun to Earth’ and produce controlled³ fusion reactions, at rates sufficient to generate more

¹ The exception of course being the fissile elements, although as they themselves were most likely created in the explosion of another star, the energy of which came from fusion reactions, the exception proves the rule.

² Although hydroelectric energy is both renewable and carbon free, the additional environmental costs which come with damming rivers and flooding valleys means the future will most likely see fewer hydroelectric plants not more.

³ It bears emphasizing that ‘controlled’ is a key qualifier here; humanity has already mastered obtaining

energy than the method consumes, using as fuel light elements abundantly available on the earth's surface.

Both approaches have their advantages and disadvantages. For the time being the primary advantage of the first method will probably win out, that being that we have essentially already developed the necessary technologies to begin implementing it. However energy production directly from the sun does have some complications as well, the most obvious being that all methods have certain environmental requirements (clear skies for solar and prevailing wind currents for turbines, for instance) which aren't necessarily available in all regions or at all times of the day. In addition solar power has a fundamental space efficiency limit of 1000 W/m^2 [4], that being roughly the intensity of solar radiation at the radius of the earth's orbit (and of course it is lower than that on the earth's surface at the higher and lower latitudes) meaning solar power plants will have to become very large installations as they are scaled up to meet energy needs. However, solar panels on rooftops can currently be used to meet the power needs of individual homes.

Unfortunately at the moment net-energy positive controlled thermonuclear fusion reactors are not merely inefficient, expensive, environmentally damaging, or dangerous, but rather are entirely non-existent. Therefore, it is somewhat difficult to speak with specificity about what a fusion power plant might look like. However it is fair to say it will not depend on certain weather or geographic features in order to function. Although current experimental and theoretical designs have tended to be rather large scale, a fusion power plant should still have a much smaller footprint than a comparable solar installation [5]; furthermore there is no theoretical lower limit on the scale of fusion (contrasting among other things with fission reactions), meaning it is not unreasonable to postulate the eventual development of 'table top' fusion energy sources. Looking much farther into the future perhaps, if human society is ever going to venture beyond this solar system, it's hard to imagine doing so powered by anything other than controlled fusion energy.

Therefore, we argue that, for these reasons among others, even if practical fusion reactor

net energy gain in uncontrolled fusion, in form of thermonuclear weapons. However no efforts to exploit that energy for power production has been successful so far. Even if that were not the case, fusion bombs always require a fission 'primary' to initiate the reaction, the use of which has many undesirable features.

designs are confined to the moderate to distant future, and in the immediate future more traditional ‘renewable’ energy technologies as well as fossil fuels are expected to dominate the global energy economy, it is nonetheless not just practical but imperative that we continue to research fusion science with the ultimate goal of controlled fusion energy.

The research presented in this dissertation will explore a piece of a potential fusion energy concept – the acceleration and transport of highly energetic electrons using high intensity lasers as part of the Fast Ignition concept [1], which is itself an extension of Inertial Confinement Fusion [6]. The primary method by which we will research this topic is through computer simulation, specifically through the use of Particle-In-Cell codes [7, 8]. As these codes were not originally designed for the parameters and scenarios which come up in this research, in doing our work we have discovered and explored certain limitations and solutions to using this method, for instance modeling plasmas where collisions are important. Where appropriate we will go into detail on these numerical methods as well. It is also worth noting that the physics of how an intense laser is absorbed at the surface of a solid material is both basic in nature and of relevance to other applications such as compact sources of 100 MeV or higher energy protons [9] as well as for exciting collisionless shocks in the laboratory [10] which is of interest in astrophysics and space physics.

1.1 Fast Ignition Inertial Confinement Fusion

Although the categorization is not fundamental and historically there have been other approaches explored, currently the vast majority of major fusion research can be said to be in one of two fields: either Magnetically- or Inertially-Confinement Fusion [11, 12].

Magnetic confinement fusion (MCF) [11] is reasonably well summarized by its name – it is an approach to fusion in which the fusion fuel is contained within the reactor vessel at the necessary temperatures and pressures (densities) primarily through the use of a strong magnetic field. The most heavily used and studied magnetic confinement device is known as the tokamak. The name is an anagram of a Russian phrase (having been first invented there), meaning something along the lines of ‘toroidal chamber with magnetic coils.’ Simplifying

somewhat, it consists of a toroidal or donut-like [13] plasma chamber surrounded by coils wrapped in the poloidal (‘through the donut-hole’) direction; by energizing these coils and at the same time driving a current in the toroidal (‘around the donut’) direction in the plasma itself, a twisting magnetic field which wraps in both the toroidal and poloidal direction at once can be established. This twisted magnetic field, and the self-consistent flow of the plasma along it, greatly enhances the confinement of the plasma as compared to the case where only the poloidal coils carry current, as it forces the particles to share their time between the outer region of the torus, where the plasma naturally drifts towards the walls, and the inner region where the drift is in the opposite sense. This is a promising candidate for a practical fusion reactor and much exciting work is currently being done in this field [14]; however we won’t discuss it further in this dissertation.

An alternative approach to fusion is what is known as Inertial Confinement Fusion (‘ICF’)[12, 6]. The author has to admit that he frankly does not find the name very illuminating to anyone who doesn’t already understand the concept. Sometimes it is colloquially referred to as laser fusion, which although inexact (ICF can also be driven by x-rays or ion beams, not just by lasers) is somewhat more descriptive. When ICF is driven by x-rays they can be generated by high energy lasers hitting a wall made of high Z material or through the compression of a separate plasma by an energetic z-pinch.

At its simplest description, ICF consists of a spherical pellet of cryogenic fusion fuel which is illuminated evenly from all directions either by x-rays, by an array of high-powered lasers, or by an array of ion beams. There are other ICF schemes which utilize a cylindrical implosion [15] but we will not discuss these further. In the case of x-rays, they can be generated by focusing many laser beams into the holes of a cm scale “can” called a hohlraum. The lasers hit the inner surface of the hohlraum which converts the laser energy into x-rays which then fill the hohlraum. It is these x-rays which drive the implosion of the fuel. This is called indirect drive ICF. The lasers can also be focused onto the fuel pellet directly. The energy of the laser beams is absorbed by the target surface, heating the surface and causing it to expand outward (this is known as ablation, and the outer layer of the target is termed the ablator). By Newton’s third law, the outward force accelerating this expanding material

must be exactly equal to a force pushing inward on the remaining target, and this inward force causes the target to implode. The implosion raises the temperature and pressure at the center of the fuel to levels sufficient to initiate fusion. When the laser beams directly ablate the surface it is called direct drive ICF and when the laser beams are converted to x-rays when they ablate the surface it is called indirect drive ICF.

As in MFE, in ICF there is a “Lawson criterion” for ignition [16], i.e., the density (n) times the confinement time (τ) needs to exceed a value which is a function of the temperature, T . This condition is smallest for temperatures between 10 and 15 keV which are possible in a tokamak. However, in ICF it is not possible to achieve such high temperatures, e.g., the temperature of the hot spot (to be defined later) are typically closer to ~ 5 keV. However, in ICF the temperature rises once ignition occurs due to the alpha particle heating.

In ICF the confinement time (inertial time) roughly scales as the radius of the fuel, R , divided by the sound speed, c_s which is proportional to $\sqrt{T/M}$ where M is the ion mass. If the Lawson criterion is written as $n\tau > c(T)$ then it follows that it can be rewritten as $\rho R \sqrt{1/TM} > C(T)$ which implies that the Lawson criterion for ICF scales as $\rho R > \sqrt{M}f(T)$ where $f(T)$ is another function of temperature. Therefore ρR is a fundamental parameter in ICF. The above arguments are rather simplistic as they do not address isobaric targets in which the fuel is compressed such that a cold dense fuel surrounds a less dense hot spot with each component in pressure balance. For such cases the radius of the entire fuel, R_{cf} determines the confinement time while the hot fuel is what first ignites. Furthermore, the temperature of the cold shell and hot fuel are different. More accurate arguments show that the confinement time actually scales as the cold fuel radius divided by the peak shell velocity, V_i [17]. However, the prominence of ρR in the Lawson criterion for ignition remains.

The combination of the density times the radius, ρR , is termed the columnar or areal⁴ density, and it can be seen to have units of mass per unit area. The areal density is also important for another reason, as when multiplied by some cross section it would approximately give the fractional rate of events. Indeed, the areal density of the hot spot must exceed the alpha particle stopping power in order for it to sustain a burn and ultimately ignite the

⁴ This is the adjective tense of ‘area.’

entire fuel. In addition, the goal is for the most of the cold fuel to burn and the burn fraction also depends on its ρR . The requirement that the stopping distance of the alpha particles be less than the hot spot radius is ρR must exceed roughly $.3g/cm^2$, the requirement that the hot spot can ignite is ρR exceed roughly $1g/cm^2$, and the requirement that most of the cold fuel burns is ρR exceed roughly $6g/cm^2$. These numbers are based on a deuterium-tritium (“DT”) fuel.

An important consequence of the requirement that ρR exceed some critical value follows from the fact that the total mass is proportional to ρR^3 . Taken together they imply that the required density scales as the inverse root of the total mass, i.e., smaller masses require higher compressions. The densities in fact turn out to be very high indeed; for 10mg of deuterium-tritium an areal density of $6g/cm^2$ implies a mass density of $300g/cm^3$, over thirteen hundred times the density of DT ice, $0.225g/cm^3$. Note that typical masses currently being used at the National Ignition Facility are even less than this, with less than a milligram of total fuel [18]; however the areal density is also significantly less, roughly $1 - 2g/cm^2$ [19].

One might be tempted from these numbers to solve the problem of needing high compression by simply using more fuel; 10 mg certainly doesn’t seem like very much. Unfortunately the very high caloric value of the fuel leads us to another constraint on the reaction – that we not inadvertently make a fusion bomb. For example, 10 mg of equimolar deuterium-tritium on full thermonuclear burn-up releases over 3GJ of energy, the equivalent of 800 kg (nominal) of trinitrotoluene. Luckily that (rather alarmist) comparison somewhat exaggerates the destructive power of the fuel, as that energy is carried by high-velocity and hence at low-momentum particles. None-the-less, although the specifics are a matter of engineering and not subject to a hard cut-off, such very large energies mean only a few 10’s of milligrams at most can be reacted at once without the risk of destroying the reactor vessel. In addition to trapping us in the high-compression regime, this limitation on total fuel has the added disadvantage of limiting the overall efficiency.

Another important point is that it isn’t in fact necessary to heat the entire target to thermonuclear temperatures, which are on the order of 10 keV independent of density. Instead the best approach is to have a small volume, known as the ‘hot spot,’ within which nuclear

fusion is self-sustaining. Since the heat needed to initiate the fusion reactions is going to be supplied by ignition of the hot spot, we are free to have the rest of the fuel be any temperature we want, and therefore we choose to have it be as cold as possible. ‘As cold as possible’ for highly compressed matter means a Fermi degenerate gas, which has an average energy per particle given by $3/5 E_F$ where E_F is the Fermi energy,

$$E_F = \frac{\hbar^2}{2m} (3\pi^2 n_e)^{2/3}$$

For $300g/cm^3$ DT the Fermi energy for the electrons is about 600eV, so the savings from having cold fuel can be substantial⁵. In reality of course there will be a deviation from perfect degeneracy; it is common to parameterize the quality of an implosion by what is called the isentrope parameter, α , which measures the ratio of the actual energy to that of a Fermi gas. Values of α of between 2 and 4 are not considered unreasonable, and so the savings can still be meaningful.

One fact that may not be at first obvious is that, for a hydrodynamically assembled fuel mass, the density of the main fuel must be much higher than in the igniting hot-spot, such that the pressure is approximately equal everywhere, that is to say the target must be approximately isobaric[20]. The reason is that, were the hot spot and the main fuel to instead be at equal densities – and therefore different pressures because of the difference in temperature – the interface between them would be the edge of a shock front, which according to the so-called jump conditions (which, in this case, comes simply from the conservation of momentum) cannot have zero velocity but rather must still be moving inward. One can either say that the system we’re describing therefore isn’t that of the finally assembled fuel mass, or simply say that the bulk kinetic energy of the main fuel is ‘wasted.’ Either way the conclusion is that the ideal ignition configuration is that of a low-density, high-temperature hot spot, surrounded by a high-density, low temperature fuel. The consequence of this is, for a given hot-spot density (or equivalently pressure, since temperature is fixed) the density of the main fuel is specified, beyond just the condition that ρR be sufficiently

⁵ For the ions the savings is of course much greater since the Fermi energy is inversely proportional to the mass; it turns out to be less than an electron-volt for almost any density imaginable. Their energy in the main fuel is therefore trivial, even at a high isentrope.

high.

Lastly, returning to the very high compression required for ICF, an important consideration is that such convergence ratios require extremely even illumination over the entire target surface. To understand this fact it is sufficient to imagine squeezing a balloon between one's hands – as you try to push it in it will respond by slipping back between your fingers. Any irregularity in force on the surface of the target can be expected to lead to irregularities in the final fuel configuration which are at least linearly proportionate to the convergence ratio, and so higher densities and compressions require ever high levels of radiation symmetry.

For direct drive there are potential solutions to generating the necessary evenness using laser optics alone. For instance one approach is to defocus the laser beams so that the spot size is an entire hemisphere [21]. This leads to much overlap between adjacent beams and so much more uniformity; a disadvantage of this approach is that as the target implodes it becomes smaller than the laser focus, and so the outside of the laser beam then misses the target and is wasted. Other techniques can also be used concurrently with wide beam focus, such as induced spatial incoherence (ISI) [22], however there is no 'magic bullet' and some degree of laser light irregularity will always be present.

In indirect drive the fuel capsule is enclosed inside a larger hollow capsule known as a hohlraum (a German word meaning not too descriptively 'hollow room' or 'empty area'). This hohlraum is made of a heavy metal, usually gold, and is generally in the form of a cylinder with two small openings at either end (termed, straightforwardly, laser entrance holes [LEH]). To compress the fuel pellet, the laser beams are focused through these openings onto the inner surface of the hohlraum, where they are absorbed. The gold then thermally re-emits the laser energy in the form of x-rays, which then bath the fuel capsule more evenly from all directions and angles (though perfect uniformity is still difficult to achieve).

Though it is far from clear that direct drive implosions will not be able to generate fusion ignition and gain, at the time of writing most effort and expense in central hot-spot ICF research is focused on the indirect approach. Indirect drive, however has a number of disadvantages which make it less than ideal. Although the coupling of the x-rays to the fuel

is in fact more efficient than that of the laser light, the same cannot be said of the coupling of the lasers to the hohlraum wall, and overall the energy cost of indirect drive is higher than direct drive. The cost and complexity of target fabrication is also greatly increased, for obvious reasons. Also, within the concept of a commercial power reactor with a high shot rep-rate, the loss of two axes of symmetry imply much greater difficulty in target alignment, to say the least.

Because of all these difficulties and constraints associated with the isobaric condition of central hot-spot ICF, if possible it might be beneficial to remove the construction of the high pressure hot-spot from the fuel assembly and so reduce the densities and compression ratios present during the implosion. This reduces or eliminates hydrodynamic instabilities which relaxes constraints on the uniformity of the implosion. One approach which has been put forward which accomplishes this task, at least in theory, is what has become known as Fast Ignition Inertial Confinement Fusion [1, 23] (‘FI ICF’ or just ‘FI.’)

The basic idea behind Fast Ignition ICF is to separate the task of compressing and igniting the fuel (the latter of which must happen much more quickly than the former, hence the name.) At its simplest, FI begins in a manner essentially similar to traditional ICF, with the compression of the fuel target using symmetric laser irradiation on an ablator surface. The difference is that the geometry of the target and temporal profile of the driver beams is such that no central hot-spot is formed. Instead at the point of fuel stagnation when the compression phase ends, all the fuel is a cold, uniformly dense mass (confusingly termed ‘isochoric,’ the etymology of which actually means equal volume), which critically can be colder and less dense than the main fuel needs to be in traditional ICF; this also means it can have a larger radius and hence a longer confinement time. At this point a second laser system is switched on – consisting of a single, narrow beam of very high intensity, 10^5 - 10^6 times the peak intensity of the compression beams – and a portion of the energy of this beam is converted into the kinetic energy of an electron beam which is created near the critical density of the plasma. The location of this critical surface is some relatively far distance from the highly compressed fuel. These electrons penetrate to the fuel and heat a small volume to thermonuclear temperatures, the radius of which again is determined by a condition on its

areal density ρR (slightly more restrictive than in the isobaric case, but comparable) which then propagates a burn wave through the rest of the fuel just as before.

Due to the nature of ablation physics, it is not the case that the assembled fuel mass can be expected to be a hard sphere of uniform density and crisp edges. Rather there will be a coronal plasma entirely surrounding the dense fuel, of approximately exponential dependence on the radius. The maximum density through which the laser is able to propagate, known as the critical density (n_c), is orders of magnitude less than that necessary for thermonuclear burn; and the point at which the plasma reaches this density, known as the critical surface, can be many 10's, even 100's, of microns away from the point of maximum compression. Heating the core via thermal diffusion from the point of laser absorption through this density gradient would be far too slow for ignition, in addition to being very energetically wasteful. Therefore rather than couple thermally to the plasma and effect an isotropic heating – as is the case for the drive beams – for the ignition beam it is instead necessary to couple kinetically directly to the electrons, creating a highly directional beam of relativistic particles. These relativistic electrons then travel, ideally, unimpeded up the density gradient before, ideally, stopping fully in a small space in the core.

Furthermore, the laser propagation and absorption processes are nonlinear and complicated. For example, the laser is not actually absorbed at the standard critical surface, that is the point where a small amplitude laser would be reflected. Rather, in the initial FI scheme non-linear effects caused by the high intensity of the ignition beam (or possibly even another, earlier beam) are used to ‘hole bore’ through the lower density coronal plasma, so that the laser energy is transferred to electrons at a highly over-critical surface closer to the core. Alternatively, a slightly modified scheme has since been proposed, in which a gold cone is embedded in the cryogenic fuel, which is then compressed around the tip of the cone. Ideally the cone itself is unaffected by the compression phase, and the ignition laser can then be fired into the cone funnel, allowing it to reach a dense surface close to the core without needing a hole-boring phase. In addition, how the laser energy is converted into electron kinetic energy is complicated and not fully understood.

It this dissertation we will examine the laser absorption and electron transport phases

of the FI scheme, especially how electrons gain energy and are sent forward into the dense plasma. We will not be investigating either electron stopping or hole boring. Based on electron stopping considerations, the original FI scheme imagined a goal of having electrons in an energy range of 1-3MeV; this energy range is found by assuming the particles will be slowed down through stochastic Coulomb collisions with the cold core, and then equating their stopping range to that of the alpha particles so that they will both be absorbed in a hot spot of the same size. It may however be the case that collective wakefield effects in the very dense core – for instance similar to the simulation effects mentioned in Chapter 4, but for physical electrons – would enhance the stopping above that predicted by collisional stopping alone; such effects would have the outcome of widening this window somewhat, making ignition slightly easier [24]. Furthermore we will be looking at absorption at an initially smooth, highly over-critical surface. The insights gained from this geometry will somewhat be able to be extended to a hole-boring experiment; however they will be most pertinent to the cone-in-target concept.

1.2 Accepted and Proposed Mechanisms for Light Absorption

Although in Chapter 2 we will be identifying and demonstrating what the main absorption mechanisms for FI (for an intense laser impinging on a sharp overdense interface at normal incidence) in is worthwhile to first review what they are not. that way we may better understand the challenges and constraints on understanding the physics involved.

The physics of the absorption of electromagnetic radiation by matter is almost impenetrably complicated, and even limiting our discussion to mechanisms pertinent to plasma physics results in a very diverse range of phenomena – a consequence among other things of the orders of magnitude spanned by the parameters of both plasma and light. Therefore we will review only some of the more commonly encountered or quoted mechanisms here.

The most basic absorption mechanism for a plasma is collisional absorption, or as it is also known, inverse bremsstrahlung⁶. Although the details of the process can be rather

⁶ ‘Bremsstrahlung’ is a German word meaning ‘braking radiation’ and refers to the light given off by a high-

complicated[25, 26, 27, 28, 29, 30], it turns out that in many cases it is accurate to simply treat it as a loss term in the plasma dispersion relation caused by an effective collision frequency[27]. In the weak field limit and for plasmas near equilibrium the effective collision frequency can be taken to be simply the classical Coulomb collision frequency (up to a geometric factor of order unity):

$$\nu_{eff} = \frac{4}{3}(2\pi)^{1/2}Z^2e^4N_iL/m^2v_{th}^3$$

where Z is the charge state of the ions, N_i is the ion number density, e is the charge of the electron and m its mass, v_{th} is the characteristic thermal velocity of the plasma distribution function, and L is Coulomb logarithm. For the strong field limit ($eE_0/m\omega_0 \gg v_{th}$) the equation has a similar form [31], although the thermal velocity is replaced by the quiver velocity, v_{os} , which is the peak velocity the particle gains oscillating in the laser electric field, $v_{os} = eE_0/m\omega$:

$$\nu_{eff} = 16Z^2e^4N_iL/m^2v_{os}^3 (\ln(v_{os}/2v_{th}) + 1).$$

The thermal velocity now enters only through the Coulomb logarithm and a new logarithmic term comparing it to the quiver velocity

There are a number of reasons inverse bremsstrahlung would not be imagined to be the acceleration mechanism for FI electrons; we highlight two. First inverse bremsstrahlung occurs as the light travels through underdense plasma, and as stated earlier in FI ideally the laser travels through vacuum or at least very low density plasma before striking a highly overdense “wall” and stopping. Second, as would be imagined, FI is very much in the strong field limit of inverse bremsstrahlung, and the quiver (proper) velocity is on the order of the speed of light or greater. Therefore the effective collision frequency is very, very low, the plasma is essentially collisionless and collisional absorption is not energetically important.

One of the first mechanisms proposed for the absorption of intense laser light is what has become known as resonant absorption[32, 33], which is an accurate if somewhat generic

energy charged particle as it slows down in the Coulomb field of another charge. Calling the absorption the ‘inverse’ of this comes from the fact that the quantum mechanical calculations of the emission and absorption of light can be formulated as the same transition probabilities with the final and initial states switched; both, however, are in fact entropy creating events in reality, and so though they may be mathematically inverse-able they are not thermodynamically reversible.

description. It is a collisionless and linear phenomenon which couples the laser energy directly to a plasma oscillation⁷. Qualitatively the ideal case is rather simple to explain. It involves a plane wave laser traveling up a linear density gradient at an angle to the direction of the gradient (obliquely incidence). Although a wave directed parallel to the gradient direction (normal incidence) would travel up to the critical density, where the laser frequency and the plasma frequency are equal, that is not the case for a wave traveling at an angle. In this case the density gradient refracts the wave so the wave vector follows a curved path which ends up perpendicular to the gradient at some finite distance away from the critical surface, after which it turns back out of the plasma. However, although the wave doesn't propagate past this turning point, the field does not go straight to zero there; rather it has an exponentially decaying profile (it classically tunnels). Therefore, there will be some finite value of the field at the critical surface, where the laser and the plasma are resonant to each other, and a Langmuir wave will be driven at the critical surface. Note that the polarization of the plane wave is key to driving the plasma wave, as the electric field must be in the plane defined by the wave k -vector and the gradient direction, so that the electric field points in the direction of the gradient at the point of resonance; if the electric field instead points perpendicular to the gradient, although it may create a current it will not create a charge separation and so there will be no plasma wave. Although in the linear regime resonant absorption can be rather efficient, with energy conversion of around 50% for optimally matched parameters, we would not expect it to come into play in the FI scheme, first, because it's not going to be in the linear regime, second because the density gradients will be much too steep so there is essentially no resonant layer, and third because the laser propagates with nearly normal incidence.

For completeness, we mention in passing another set of non-linear couplings between light and plasma, though as with collisional absorption these only occur in underdense regions and so will not be a (desirable) part of the ignition phase of FI. As explained below, they probably *will* be a part of the compression phase of FI, as they are in traditional ICF;

⁷ Although note that the way the plasma oscillation is converted to energetic electrons is related to what is often called wave breaking, which is nonlinear.

however again they are undesirable. These are what are known as parametric instabilities. A parametric oscillator is one governed by the following equation[34] (in the case of no damping and no driving):

$$\ddot{x} + (\omega_0^2 + k(t))x \equiv \ddot{x} + (\omega_0 + \Delta\omega(t))^2 x = 0.$$

One way to interpret this equation is that the natural frequency of the oscillator, which is the one parameter of the equation, is a function of time, hence the name *parametric instability*. Without loss of generality, provided the time-dependent parameter remains bounded, we can assume it has vanishing time average, $\langle \Delta\omega(t) \rangle = 0$, and so it is usual to think of it as defined by some harmonic series. The important point about parametric oscillators, for the present case, is that if there is some harmonic relationship between the natural frequency of the oscillator (ω_0) and the frequency of the perturbation, the oscillator can be unstable – even in the absence of an explicit driving force, the amplitude of a parametric oscillator can grow exponentially in time, provided there is some small displacement in the system from the zero (equilibrium) state; this contrasts with the case of a driven harmonic oscillator for which the growth is linear in time (and a seed fluctuation is not necessary). The pertinence of this equation to a laser-plasma interaction arises in that when a laser propagates in an underdense plasma the non-linear forces associated with that light can modify some parameter of the plasma – for instance the density – and if that modification is near resonance with another mode in the plasma, that mode can grow exponentially in time, seeded by the plasma’s thermal fluctuations. The actual set of equations are, however, more complicated than the parametric oscillator.

Because these excited oscillations are also waves, an important further restriction come from what is known as classical frequency and wavenumber matching. These are related to energy and momentum conservation, which can be appreciated by multiplying the equations by $\hbar$. Therefore, for three wave processes, the sum of the frequencies and wavenumbers of the excited waves, often called daughter waves, have to equal those of the original wave

$$\omega = \Sigma \omega_n$$

$$\vec{k} = \Sigma \vec{k}_n$$

In an unmagnetized plasma there are three types of fundamental oscillations possible: electromagnetic waves (‘light waves’), electrostatic waves (‘plasma waves’), and ion-acoustic waves (‘sound waves’). Their dispersion relations are such that the frequency and wavenumber matching are only satisfied for an electromagnetic wave if there are at least two daughter waves, and at least of one those daughter waves is not also a light wave. Otherwise most combinations are possible, and each is its own named phenomenon: the decay of a light wave into a plasma and light wave is known as Raman scattering, the decay of a light wave into a sound and a light wave is known as Brillouin scattering⁸, the decay of a light wave into a plasma and a sound wave is known (somewhat generically) as parametric decay, the decay of a light wave into two plasma waves is known as two plasmon decay, et cetera. There are other instabilities that involve four waves [35] (often called modulational like instabilities) as well as multi-dimensional effects such as side scattering.

As can be inferred from all these names, parametric instabilities are rich in physics and heavily studied phenomena [36, 37]. It is not referred to as the nonlinear optics of plasmas. The ICF community is especially interested in them, as their presence during the implosion phase can greatly undermine compression quality, by scattering energy back from the target, by scattering energy onto the target in an uneven way, or by generating energetic electrons. However again they are certainly not the mechanism for electron acceleration and laser absorption at a sharp surface, as is needed for the ignition phase, and furthermore they still wouldn’t be expected to occur even if there were significant underdense plasma which the ignition beam had to travel through, as the beam is actually too strong for these mechanisms to have an appreciable effect on the energy balance. Therefore we will not discuss them further.

Next, we arrive at a pair of absorption mechanisms for which it is not *a priori* obvious that they would not be appropriate for FI parameters, and hence they are sometimes invoked as absorption mechanisms for intense light. They are known as ‘Not-so-resonant’ resonant

⁸ Both Raman and Brillouin scattering in a plasma are named in analogy with phenomena from non-linear solid state optics. Personally we feel that, although the two Brillouin effects are comparable, the two Raman effects really are distinct, as in the solid state case it is a purely quantum mechanical effect. This terminology is well established, however.

absorption and $J \times B$ heating. Although these are distinct mechanisms, both theoretically and phenomenologically, we group them together because they do occur under somewhat similar conditions, and so are often seen together in the literature [1, 38, 39, 40, 41, 42, 43]. However, both also miss a facet which we will find to be key in our absorption simulations, that the absorption can be greatly affected by the temperature of the target plasma.

First let us discuss ‘not-so-resonant’ resonant absorption [44]. As this is a bit of a mouthful, it is often referred to as either Brunel absorption (after the original author) or rather vaguely as vacuum heating; we prefer the former. As can be inferred from the full name, it is related to the resonance absorption described earlier, specifically in that both involve the electric field of an obliquely incident laser coupling efficiently to the plasma electrons as it turns normal to the surface at the plane of reflection. The difference arises in the fact that, as Brunel observed, for very intense lasers the profile of the plasma would be compressed up to a sharp surface, so that the density abruptly rises from underdense to greatly overdense and hence nowhere is the plasma resonant with the laser frequency (no plasma wave can be excited). In such a geometry, the laser will also be reflected abruptly, and in the limit of full reflection the electric field at the surface will be normal to the surface and will have a peak electric field twice that of the vacuum field. This field is strong enough that it may pull electrons out of the plasma into the vacuum region before pushing them back in; this movement is an-harmonic. The electrons feel the laser field as well as the fields from other electrons that have been pulled into the vacuum. Brunel showed this leads to net energy gain and absorption during the half cycle for which the electric field pulls electrons into the vacuum, then during the other half cycle they are pushed back into the plasma.

We do expect Brunel absorption to play a part in the absorption for certain situations and geometries; indeed we believe it may account for some of the late-time heating we describe in Chapter 2. However it involves a few features and approximations which will make it inconsistent with our findings for the main absorption mechanism. The most important of these is that Brunel absorption goes to zero at normal incidence, but we find there can be non-trivial absorption at normal incidence in both 1D and higher dimensions. In addition, it is based on a dipole approximation for the laser, it does not take into account the laser

magnetic field, it predicts absorption every laser cycle, and, as mentioned earlier, it does not depend on the plasma temperature. All of these facts are in conflict with our results.

Next, we come to so called $J \times B$ heating. The original, rather terse paper [45] derives this effect using the ponderomotive force arising from the gradient in light intensity as a laser is reflected over a thin volume in front of an overdense target. They show that the force, thus derived, is composed of a constantly forward-pointing component and an oscillatory component with a frequency twice the laser frequency. They then equate this oscillatory force with an electrostatic field, E_d , near the surface and within the skin depth

$$\frac{eE_d}{m} = \frac{1}{4} \frac{\partial}{\partial x} v_L^2(x) (\cos 2\omega_0 t)$$

where v_L is the local quiver velocity, that is the peak strength of the field normalized to a velocity $v_L = eE_L/m\omega_0$. They propose that this oscillating field can lead to heating of the plasma, just as in the case of Brunel heating, while the constant force compresses the plasma to a strong gradient.

This paper has unfortunately caused a lot of confusion, in our opinion⁹. The brevity of the explanation, coupled with the name $J \times B$ – meaning, essentially, little more than that the magnetic field is involved somehow – has left a lot of room for individuals to interpret what exactly is meant. However, the magnetic field cannot do work, or give energy to the particles. It is implied in the original paper that the work is done by a longitudinal E-field driven by the $J \times B$ force. In addition, their derivation is based on the ponderomotive force approximation (which they derive ‘in the usual way’) which assumes that the scale-length of the field gradient is much greater than the typical distance traveled by the electron in each oscillation. This is explicitly not true in our case, as the plasma surface has been compressed. The term $J \times B$ is often invoked any time there is a response in the plasma at twice the laser frequency, but as we describe in Chapter 2, there are several acceleration processes at twice the laser frequency.

Therefore it is somewhat hard to prove that a given interaction is *not* $J \times B$. However we find a few reasons to believe this is not the main absorption mechanism for FI relevant laser

⁹ This confusion is not the fault of the original authors, but rather the paper being referenced to explain any process in which electrons are sent forward twice each laser period.

beams. First, the absorption is expected to be a decreasing function of plasma density, but we find significant absorption with densities 100 times the critical density (though note that for *cold* plasmas our simulations in fact confirm this prediction.) Second, again as mentioned earlier, there is no account taken of the effects of the plasma temperature. The very fact that the plasma temperature affects the absorption is related to the fact that the motion and forces experienced by the electrons are more complicated than the first order ponderomotive approximation would imply.

The (one dimensional) simulations described in the original paper on $J \times B$ heating also use laser intensities a few orders of magnitude less than those involved in FI, and so it might be that this mechanism is valid under those conditions and not under ours. Alternately their description may be an approximation of the mechanism we propose, as they do see some dependence on temperature which is consistent with our findings. We have not had an opportunity to revisit their parameter regime in depth and so cannot say for certain what exactly the authors were seeing. But we assert that it is not consistent with the absorption we see in our higher intensity and FI relevant simulations.

There are a number of other absorption mechanisms at sharp, flat surfaces which have been proposed recently, often specifically with regards to the FI concept (e.g. [46, 47, 48].) Some common themes in these mechanisms are that absorption can occur in one-dimensional simulations, absorption can occur for cold plasmas, the longitudinal or ‘electrostatic’ field created by charge separation at the surface is important, and a lack of an inclusion of the magnetic field at the surface when analyzing the particle motion. We will show that we find all of the features to be inconsistent with what we observe in our simulations.

1.3 Particle-in-Cell Codes and OSIRIS

In general for a group of n charged particles interacting through electrostatic forces, there are $\mathcal{O}(n^2)$ interactions happening at any given time, since each particle is interacting with all the others. Simulating such a group of particles therefore quickly becomes unfeasible as n becomes large. If the fields are extended to electromagnetic interactions with a finite speed

of light, then a further complication arises in having to calculate the forces on particles at a distance based not only on their current positions but also on all of the positions they had at some time in the past.

The Particle-in-Cell or PIC algorithm[8] solves both of these issues, by having particles interact not with each other directly but rather through a discretized field object, which reduces the complexity and the number of floating point operations significantly. The particle advance will scale as $\mathcal{O}(n_p)$ where n_p is the number of particles, while the field solver will scale as $\mathcal{O}(n_g \log(n_g))$ where n_g is the number of grid points for a Poisson solver and as $\mathcal{O}(n_g)$ for a finite difference Maxwell solver. A cartoon summary of the PIC ‘loop’ is shown in figure 1.1. The basic idea is summarized thusly: First one defines some region of space (in any number of dimensions). Within this space there exist a finite number of charged particles whose positions are allowed to exist at any point¹⁰ in this continuous space. It is useful to view the particles themselves not as point particles but rather as ‘clouds’ of charge, with a finite size and a specific rigid shape. Each particle also has associated with it a velocity, or momentum, which again can take on any continuous value. This velocity is then used to advance each particle’s position in space through some finite time-step (the same for all particles.) Then next, this motion – along with the particle’s charge – is used to calculate a current associated with each particle. Since the particles are not delta functions, the current and charge are also volume currents and charges. However the charge and current are not deposited in a continuous space, but instead are weighted to a discrete set of points on a grid near the particle. Usually this grid is regularly distributed throughout the whole volume of the simulation. The grid is needed to solve the field equations. For the research in this dissertation we use the full set of Maxwell’s equations for the field equations, though the time independent equations (Gauss’ law and the vanishing divergence of the magnetic field) are only explicitly solved for at $t = 0$, after which they remain valid through the self-consistency of the four equations. These can be solved using finite difference approximations for the derivatives or using Fast Fourier Transforms (FFT’s).

The current is used to calculate the change in the values of the fields, through the time

¹⁰ Up to the limitations of the machine representation of floating point numbers.




FIGURE 1.1: *PIC Loop*

dependent Maxwell's equations (Ampere's law and Faraday's law). The values of the fields used on the right-hand-side of these equations will either be values calculated at the previous time step if this is not the first time step, or will have been specified by the user as an initial condition. The fields themselves are defined on a grid similar to that which contains the current object (though the exact locations of the currents and fields, or even their various components, do not in general coincide in order that the finite difference operators are centered properly). Finally the values of these fields on the grid are weighted back to the location of each particle and used to update its momentum to a new value, returning us to the beginning of the loop. This naturally follows from the idea that the particles have a finite size and shape.

This basic framework is the same for all PIC codes, of which there are many. Throughout this dissertation we will specifically be using the PIC code OSIRIS [2, 49], developed by a consortium between UCLA and IST in Lisbon, Portugal. There are a number of features

which make OSIRIS a particularly powerful and efficient code, both for the user and for the developer.

One of the most important features of OSIRIS is its efficient implementation of a hierarchy of parallelization. At the top level MPI is used between nodes with an arbitrary dimensionality for the domain decomposition. MPI can also be used across all cores on a node as well. OSIRIS has been run successfully on millions of processors demonstrating high parallel scalability on both weak and strong scalings. In addition many parts of the code, including all the main parts of the code which have significant computational load – current deposit, particle advance, et cetera – now incorporate the a second level of parallelization through the use of the OpenMP multithreading technology, allowing multiple processors to share a single task on a single node with shared memory. This can significantly reduce the need to subdivide a given problem into discrete, independent MPI tasks, which can create communication inefficiency and overhead as the domain size for each MPI region becomes too small. Lastly, all the computationally heavy parts of the code have also been extensively translated to processor intrinsics to take advantage of a third level of parallelization based on SIMD vectorization which is available on many modern processors. This vector optimization may ‘only’ give a speed-up of a factor of two-to-four on most systems, but as it comes at literally no cost for the user (and little for the developer of other parts of the code) it meaningfully extends and enhances the OSIRIS capability. We also note that recently, much work has also been done recently to port the main OSIRIS code to very highly vectorized environments such as Nvidia GPUs or Intel Phi MICs. This has already been shown to increase the speed of computation by orders of magnitude for many problems. However, unfortunately we were not able to take advantage of this for any of the simulations performed in support of this dissertation, and so we only mention it for completeness.

Another strength of OSIRIS compared to many other PIC codes is the implementation of higher-order particle shapes – that is the specific method for depositing charge to the grid and then interpolating the force back to the particles. The most important utility of these higher order shapes is in reducing the self-heating of the simulated plasma, and hence improving energy conservation. For codes without higher order particle shapes it is often necessary to

have the grid size resolve the Debye length of the plasma; otherwise the plasma will self-heat until the Debye length becomes long enough that it is resolved. PIC codes which use the same particle shapes to deposit current and interpolate the force are not energy conserving. This self-heating is due to aliasing that arises because the particles have continuous positions and therefore can have currents with arbitrarily large wavenumbers. These show up as small wavenumbers on the grid due to aliasing. Higher order particle shapes reduce aliasing. By using these smoother shapes, along with filtering of the currents and potentially the fields, OSIRIS is able to run with much larger cell sizes or cooler plasmas (or both). Indeed for most of the simulations used for this dissertation the scale length which needs to be resolved is not the Debye length but rather the skin depth, which for non-relativistic plasmas can be many orders of magnitude larger. This greatly enhances the speed and efficiency with which we are able to run these simulations.

For the developer by far the most important aspect of OSIRIS is its adherence (roughly) to the Object-Oriented programming paradigm. In short, this means that one part of the code can operate and be modified without having to worry or even know what other parts of the code are doing (including the parallelization). In the context of our research, this meant adding things like particle splitting, advanced boundary conditions, and most especially collisions – in addition ‘non-physics’ tasks like diagnostics or set-up routines – could be done with only minimal modification to the main code. This made modifications quick to add to the program, easy to compartmentalize when troubleshooting or benchmarking, and (with caveats) easy to share with others.

1.4 Collisions in PIC

At first it might seem surprising to hear it asserted that Particle-in-Cell codes are most naturally suited for studying collisionless phenomena, that is those in which the effects of the discrete particles can be ignored. However, the reasoning is actually related to the nature of collisions themselves. In plasmas there is a minimum scale length (maximum wave number) that needs to be resolved in order to describe collective phenomena that involve electron

dynamics, which is Debye length v_{th}/ω_p . which then sets the size of a simulation cell. Within a Debye length the physics is dominated by collisions and the effects of collisions scale as $1/N_D$ where N_D is the number of particles within a Debye cube. In the PIC method the physics within a Debye length is minimized through the use of finite size particles. For some problems such as those being studied in this dissertation the minimum scale length for the collective modes is the skin depth, $c\omega_p$. As noted earlier if cell size is larger than the Debye length then aliasing between the particle quantities (given with continuous positions) and the grid quantities (fields) leads to the grid instability which causes the numerical heating of the plasma. We get around this through the use of higher order particle shapes, however the use of large cells can modify the numerical dispersion relation of collective modes so care is required.

The PIC method can be most properly viewed as modeling not the Vlasov equation (sometimes referred to as the collisionless Boltzman equation) but rather the Klimontovich equation for finite size particles. Each simulation represents a single realization of a collection of particles with all of their initial locations in phase space being known rather than an ensemble average over many realizations. A collision operator can be derived for this model using the same methods as for point particles. The result is a modified collision operator which smooths out interactions for impact parameters (wave numbers) less than the particle size (larger than the inverse of the particle size). Therefore, fewer particles per Debye cube are needed to keep the collision frequency much less than the natural frequency of modes of interest. The computational advantage of PIC codes lies in their ability to use many fewer electrons and ions than are present in reality to study collisionless phenomena; therefore, many real particles are represented by each simulated particle, which is hence often termed a ‘macro-particle’. For many problems, the goal is to use enough particles such that the collisionality is small enough so that collective effects can be studied. In this way the PIC algorithm can be an numerically effective technique for studying the Vlasov equation as the number of particles per cell that are needed in the simulation can be substantially less than the number of momentum cells needed per real space cell. However, the PIC method can also be used to study how a collision operator (for finite size particles) is modified when one

moves from a plasma near equilibrium to a plasma in a highly turbulent state.

The PIC method is now often used to study effects that involve the interaction of collective modes and collisions as it can be more computationally efficient to studying the Vlasov-Fokker-Planck (VFP) equation (the collision operator is in the form of a Fokker-Planck operator). In this case, the goal is to have a collision operator within the PIC algorithm that accurately models Landau-Boltzman collision operator or Coulomb binary collisions. We stress that the Landau-Boltzman collision operator is itself an approximation for collisions in a plasma. One's very first thought may be that we could retain just enough of the finite size particle 'PIC collisionality' to match key aspects of the "real" collisions; however as noted above while one could match the collision frequency, it is not possible to match all aspects of the collision operator in velocity (momentum) space with the necessary statistics.

The common approach to “accurately” modeling collisions into PIC codes is to increase the number of finite particles per simulation cell and Debye cell, as well as use a “fatter” particle shape to reduce the collisions between finite size particles and then add a Monte Carlo collision algorithm. In broad strokes, such an algorithm would generally start by organizing the particles into disjoint subsets of physical space, called the collision cells, such that particles can only interact with others within the same cell. The collision cell could be the same cell used for the field discretization, but in principal it could be larger or smaller. Certain bulk properties which affect the collision rate are then measured in each collision cell using all the particles, for instance temperature and density. After this, all the particles in the cell are randomly assigned into pairs, and then for each pair, using a combination of the bulk properties calculated earlier and the individual properties of the particles in the pair (most commonly just the relative velocity), a calculation is made to determine whether to collide the particles and what the details of that collision should be. The particles then exchange momentum (if a collision happens at all) in the center-of-momentum frame¹¹ so that both momentum and energy are conserved in the interaction. In most cases the decision of the specifics of the collision event is not made deterministically from these independent

¹¹ Note: NOT the center-of-mass frame, as is often erroneously stated in the literature. The center-of-mass frame is a particular choice of position variables, such that mass is centered around the origin; collisions happen entirely in momentum space, the choice of position space has no bearing on them.

variables, but also involves the use of one or more random numbers, so that the event has a statistical element to it. The use of random numbers to solve problems which may in fact be deterministic in nature (though in this case only partially so) is known as a Monte Carlo method¹².

The challenge for including collisional effects in a PIC code is hence reduced to finding the correct form for this binary collision operator that is consistent with the Landau-Boltzman collision operator. To explore the details of this collision operator, it is of course necessary to look at the nature of the underlying collision physics. It is illustrative to start with a problem which is in some ways simpler but in others more complicated, that is the solution of the Boltzmann equation for a species with an arbitrary collision law (that is more general than the inverse square Coulomb interaction) which is the subject of Bird 1970 [50], one of the earliest papers on using the Monte Carlo method to include collisions in a particle based (not a particle-in-cell) code. The Boltzmann equation, which in its generic form applies to a general gas as well as a plasma can be written in the form

$$\frac{\partial f}{\partial t} + \dot{\mathbf{x}} \cdot \frac{\partial f}{\partial \mathbf{x}} + \dot{\mathbf{u}} \cdot \frac{\partial f}{\partial \mathbf{u}} = \left(\frac{\partial f}{\partial t} \right)_{\text{collisions}} \quad (1.1)$$

where the left-hand side states that, in the absence of collisions, the convective derivative in phase space of the single particle distribution function f vanishes. This means that the phase space density of the single particle distribution function does not change as one moves with the phase space "fluid", viewed another way, this means that the volume in phase space cannot change. For a plasma the acceleration is comes from electromagnetic fields (the Lorentz force), both the self-consistently generated by the plasma particles and any applied fields. The presence of collisions (the right hand side) means that the the phase space volume can change. In essence when one combines PIC with a Monte Carlo collision algorithm it is solving the the homogeneous equation (with right-hand side equal to 0) through the PIC method with a sufficient numbers of particles to keep the effects of finite size particles sufficiently small and then including a collision algorithm which can be non-trivial.

¹² Rather unreassuringly named after the famous gambling casino - we certainly hope that our simulations are a better bet than roulette!

In general, the collision term can be expressed as [50]

$$\left(\frac{\partial f_i}{\partial t}\right)_{\text{collisions}} = \sum_{j=1}^N \int \mathbf{d}^3\mathbf{p}_j \int d\Omega g \frac{d\sigma}{d\Omega} (f'_i f'_j - f_i f_j) \quad (1.2)$$

where the sum is over all species (including species i), g is the relative velocity between the regions in velocity space between the two species, $d\sigma/d\Omega$ is the differential cross section (understood to be a function of both momenta), the prime indicates the distribution functions after the collision (that is $f'(\mathbf{p})$ in the distribution function scattering *into* $\mathbf{p}$ {this assumes reversible collision physics}), and the integrals are understood to be over all momentum space of the collision partner and over all solid angle, respectively. The reader may not be surprised to hear this is not in general an easy equation to solve analytically.

To solve the Boltzman equation numerically, Bird first assumed a form for the differential cross section, specifically an inverse power law; for instance the collision between particles in a gas can be well modeled by a force which falls off with distance with a sufficiently high power (perhaps 6 or greater). He then integrated over all solid angles to get an effective collision cross section. This cross section can then be used with the densities of the species in the simulation to calculate a collision rate for each collision cell (as a function of the momenta of the individual particles). The code would then go through each of the randomly assigned particle pairs, and would use this collision rate, along with a counter assigned to each cell, to determine whether a collision event should happen; if it should the momenta would be updated appropriately from the interaction law (the counter would be updated as well.)

Note that in this algorithm, the integral over momentum of the collision partner, $\int \mathbf{d}^3\mathbf{p}_j$ is not explicitly performed; rather by selecting two particles to collide randomly, each is sampling the distribution function of the other species in momentum space. Provided most of these pairings are ‘typical,’ the actual integral is well approximated by one or perhaps a few such samplings. The reason for choosing such a Monte Carlo method for solving an explicit integral are two-fold. First, it reduces the complexity from a many-particle interaction to a series of two-particle interactions. Second, it explicitly symmetrizes the ‘test’ and ‘field’ particles (that is species i and j in equation 1.2, respectively,) which is of course not a

physical distinction to begin with.

For a plasma interacting through a Coulomb force, this approach is neither sufficient nor necessary. To see why it's not sufficient, we observe that for an inverse-square law like a Coulomb force the integral of the differential cross section diverges logarithmically for two isolated particles. A point or finite size particle having an infinite cross section is not unphysical, it arises because the scattering force decreases with distance so slowly that a meaningful interaction takes place for impact parameters at an infinite distance. The customary way to get around this is to assert that in a plasma, where all of the particles are interacting with all the others at all times that the Coulomb force from one particle is shielded out within a distance taken to be the Debye length λ_D . The logarithm diverges not only because of large impact parameters but also because of small impact parameters. Therefore, a lower bound for the impact parameter needs to be specified. This is sometimes taken to be the distance of a 90° deflection which is the distance of classical closest approach. For higher temperatures (actually not that high) the smallest impact parameter is given by the the de Broglie wavelength. The critical point to this argument is that because the cut-off distances enter inside of a logarithm, the result is relatively insensitive to their exact values; even a change by a factor of 2 for either cut-off will have only a small numerical effect in most cases. Such a derivation is said to have *logarithmic accuracy*.

With these physical arguments, it is possible to calculate an effective collision frequency for a Coulomb interactions in a plasma; the result is (for electrons colliding with a stationary ion species at non-relativistic speeds, in Gaussian units) [51]

$$1/\tau_{coll} = \nu_{coll} = \frac{4\pi e^4 Z^2 n_i \ln \Lambda}{m_e^2 g^3} \quad (1.3)$$

where $\log \Lambda$ is the logarithmic term, known as the Coulomb logarithm; typical values are 5-20.

At this point it would be possible to use this effective collision frequency in an algorithm like the one described by Bird, and collide the particles every τ time-steps. As alluded to earlier, this is actually not necessary nor even correct. The reason for this – qualitatively and simplifying somewhat – is that typically in a plasma there are many particles within a

Debye sphere¹³, and almost all of these are farther away than the minimum cut-off distance of closest approach. These “distant” particles, which each give a small angle scattering therefore dominate the collision integral. Therefore, it is not the case that diffusion in a plasma is caused by individual strong scattering events (as in a gas) but rather by the cumulative effect of many small, and uncorrelated, interactions. Using this observation we can make another simplifying assumption, that collisions in a plasma come about due to a series of nearly continuous infinitesimal deflections; one might say these are not ‘collisions’ at all, but rather a diffusion in momentum space.

Returning for a moment to the analytical treatment, now with the assumption that small-angle scattering dominates we can revisit to the Boltzmann collision term and assume the $\Delta\mathbf{v}$ is small by making a Taylor expansion in f'_i about f_i in powers of $\Delta\mathbf{v}$. We keep up to second powers in the components of $\Delta\mathbf{v}$. This is equivalent to the Fokker-Planck expansion. After much algebra this leads to the collision term originally derived by Landau [52, 53],

$$\begin{aligned} \left. \frac{\partial f_i}{\partial t} \right|_{\text{collisions}} = & \sum_{j=1}^N \frac{2\pi e_i^2 e_j^2 \ln \Lambda}{m_i} \frac{\partial}{\partial \mathbf{v}_i} \cdot \int \mathbf{d}^3 \mathbf{v}_j \left(\frac{\mathbf{I} g^2 - \mathbf{g} \mathbf{g}}{g^3} \right) \cdot \\ & \left[\frac{f_j}{m_i} \frac{\partial f_i(\mathbf{v}'_i)}{\partial \mathbf{v}'_i} \bigg|_{\mathbf{v}'_i=\mathbf{v}_i} - \frac{f_i}{m_j} \frac{\partial f_j(\mathbf{v}'_j)}{\partial \mathbf{v}'_j} \bigg|_{\mathbf{v}'_j=\mathbf{v}_j} \right] \end{aligned} \quad (1.4)$$

This integro-differential expression has the same form as the Fokker-Planck equation, and the total Boltzmann equation with this as the right-hand side is usually called the Vlasov-Fokker-Planck (VFP) equation.

It bears emphasizing that although the underlying physics is entropy creating and hence stochastic, the VFP is a fully deterministic equation in momentum space. And, although the equation may appear formidable, in some situations – sometimes with some appropriate expansions (for instance that the distribution function is not too far from Maxwellian) – it is possible to find exact analytic solutions. In some more general cases it is possible to solve this equation directly numerically, by rewriting it somewhat so that distribution function is

¹³ In fact it is common to define the count of such particles as the *plasma parameter* and to assert that one definition of a plasma is that the plasma parameter is much greater than unity; in that sense this assertion is tautological.

expanded in spherical harmonics, a technique which was originally derived by Rosenbluth [54]. We state all this again to draw attention to the fact that although we will be using a Monte Carlo method to solve this using a PIC code, we do this with the goal to solve a deterministic equation using a computational sampling method in a more numerically efficient manner. However, it is important to recall that the VFP is actually based on the single particle distribution function which is the ensemble average of an infinite number of initial Klimontovich distribution functions. The collision term is then an ensemble average of short range variations within a Debye length. Therefore, it is not surprising that when combining PIC with an algorithm that correctly models collisions between particles within a skin depth it is necessary that this algorithm uses a Monte Carlo method. It is also worth noting that the PIC method itself has some Monte Carlo characteristics. To initialize the plasma, the velocities and hence positions are selected using a random number generator. So in order to obtain some quantities from a PIC simulation it may be best to run the same problem using different random number generator seeds and then average the results. It would therefore also be interesting to run many simulations with PIC combined with a Monte Carlo collisions with the same initial conditions for each macro-particle particle and then ensemble average the results as well as many simulations with different random initial conditions for each particle.

To solve VFP equation for plasmas in the context of a particle code, it is important to develop the correct Monte Carlo method. We start with the physical observation that motivated it: that (real) particles in a plasma undergo many small angle collisions. With this in mind, we return to a method used by Bird which randomly collided pairs of particles within a collision cell. We would then find the collision frequency for that pair using equation 1.3; and compare that frequency to the elapsed time since the last collision event. However, for a plasma because the concept of a single collision event is ambiguous (collisions occur gradually and there are many simultaneous collisions) we do not keep a counter since the last collision event; rather each particle collides as pairs at *every* collision time, and so time since the last collision event is always the sampling time, Δt_c . Consequently, we use the compare Δt_c to τ_{coll} to determine *how much* to collide each pair. Note that equation 1.4 describes a

diffusion, that is to say a point in momentum space evolves to a continuous volume. Since this is impossible in a PIC code without some form of particle splitting in momentum space (which defeats the computational advantages of a PIC code making it undesirable), we instead assign a new final momentum randomly for each, with a probability distribution equal to the actual diffused momentum. Within each pair momentum and energy are conserved. Thus, in order to include collisions consistent with physical parameters we need to add this additional Monte Carlo step to the PIC method.

This is the collision method originally developed for PIC codes by Takizuka and Abe [55] (T&A), which is the essence of all the collision methods we have in OSIRIS (and one of them being precisely the T&A method). For each pair of particles, we transform to the center-of-momentum (COM) frame, and rotate the momentum coordinates such that the particles are traveling towards the zenith and nadir. We then rotate the momenta such their new polar angle is given by θ , where θ is calculated from a second variable δ such that $\tan(\theta/2) = \delta$, where δ is a random number with a Gaussian distribution with zero mean and variance given by

$$\langle \delta^2 \rangle = (\nu_{coll}/2)\Delta t_c \quad (1.5)$$

where recall that Δt_c the time between collision calculations within the algorithm, and the factor of 2 is not well explained by T&A at all, but we can attest that does give the correct result. The new azimuthal angle is evenly distributed between 0 and 2π . When calculating $\nu_{coll} \equiv \tau_{coll}$ for each collision pair, the relative velocity is unique to each pair, the Coulomb collision comes from temperature of all the species in the cell, and the density is the lesser density in the collision cell of the two species which are colliding (for unlike particles); more will be said on this last point in a moment.

Using this method to obtain an approximate solution to eq. 1.4 necessitates that $\sqrt{\langle \delta^2 \rangle}$ is small. Otherwise, during a single collision time there will be many large angle scattering events. There must therefore be many collision events per collision time, i.e., $\nu_{coll}\Delta t_c \ll 1$. Note that this is a separate condition from the requirement that we have many pairings to adequately sample the distribution function, and in fact it is a stronger one. If we don't satisfy the latter all we do is to add a noise term to the calculation to the collision frequency,

which can still be smoothed out by increasing the number of total pairs or running many simulations and then ensemble averaging the results. If we do satisfy the former the actual nature of the collisions changes. Takizuka and Abe partially addressed this by having a cut-off for $\nu_{coll}\Delta t_c$, above which the polar angle is simply evenly distributed between 0 and π . However this is not fully satisfactory, as it is ad hoc and hence not correct, and it gives an abrupt change in behavior.

Nambu [56] realized that the cumulative polar angle deflection from many small angle scatterings is a well-defined quantity, and so it is not actually necessary to keep track of (resolve) each small angle event. Using an ensemble of simulations, he computed¹⁴ an equation which gives the correct limiting behavior of small- and large-angle scattering, and fits well the transition between them¹⁵. This method for calculating the collision angle is in all ways superior to that given by Takizuka and Abe. If $\nu_{coll}\Delta t_c \ll 1$ and there enough pairings to adequately measure the distribution function, it gives the same answer as T&A, while as $\nu_{coll}\Delta t_c$ approaches unity it still gives the correct result. Nambu’s contribution is very important. Typically, the time step used for the standard PIC part is chosen to resolve the collective modes, which for overdense plasmas scales as $\omega - p^{-1}$. For “collisional” plasmas where $\nu_{coll} \sim \omega_p$ then τ_{coll} can be comparable to the PIC time step. Therefore, in order to resolve the the collision, the collision time step Δt_c would be much smaller than the PIC times step. Therefore, with some modifications described shortly, this is the method the authors use almost exclusively with OSIRIS.

So we have somewhat come full circle. We started with the Boltzmann equation for a gas, in which ‘collisions’ look like actual collisions, i.e., scattering of billiard balls. Then, observing that this is not an accurate way to portray Coulomb collisions, we modified our numerical approach so that each collision event was a small deflection, consistent with reality. Next, by realizing that in a plasma it is not actually required that we explicitly calculate each of

¹⁴ Although Nambu’s method for calculating the collision angle is a great improvement on the small-angle formulation, and its accuracy is arguably as good as the logarithmic accuracy of the underlying physics model, it is simply not the case that they *derived* their equation, as they claim.

¹⁵ It becomes obvious when one considers it, but it should be emphasized explicitly, that in cases where the collision frequency is changing due to the collisions themselves – for instance a cold beam of electrons stopping in a background of ions – this method of taking large collision steps breaks down, and one should keep enough iterations such that the whole process is integrated decently.

those small deflections, we return to the possibility that each simulated event represented many real events, and so cause a significant change in momentum. However, we have obtained some insights through this circumnavigation. First, we are still comparing our results against the Landau collision term, eq. 1.4, therefore, the Monte-Carlo algorithm itself can be tested alone against known results. Second, the exact form of the Nanbu equation is motivated by the math of small-angle scattering, which was not obtainable from the Boltzmann collision term. Third, despite pairing particles, it is still not the case that we are modeling collision *events*, but rather just statistically relevant collision *times*; as such we are free to choose our numerical collision step as we see fit, and are not encumbered by having to keep track of any sort of collision counter.

Now is a good moment to pause and emphasize another point. Our PIC collision model is method for evolving a distribution function as predicted by a collision operator of the Landau form. However, the Landau (or equivalently the Landau-Boltzmann) collision operator is itself not exact, but rather it is accurate under certain approximations. Specifically, it is valid when the plasma parameter (number of particles per Debye cube) is much larger than unity so that there are many particles “interacting” at once. There are other methods used to derive a collision operator for a plasma such as the Lenard-Balescu method. This method views collisions as the interaction of “dressed” particles in a plasma. This approach also reveals the importance of the wakes made by test particles and the challenges posed for deriving a collision operator in the presence of strong turbulence. It can be shown that the Lenard-Balescu and Landau operators are the same when the plasma parameter is large. If the implicit or explicit assumptions used to derive the Landau operator break down it becomes inaccurate. In this case, there is absolutely no hope that this model can be used to calculate those inaccuracies, no matter how high a resolution one uses in any of the parameters – because it assumes the Landau operator is accurate from the beginning. To obtain corrections to the collision operator – which could then be fed back into a PIC collision algorithm – an entirely different approach would be needed. These include a combination of PIC and molecular dynamics where PIC is used for particles farther away than a Debye length and exact Coulomb interactions are used for those particles within a Debye length of

each other. This will not work when full electromagnetic interactions are occurring. Another option for gaining insight into collisions when the plasma parameter is small or when there is strong turbulence is to run a PIC code with progressively smaller cell sizes (much smaller than a Debye length) and with each macro-particle representing a real electron. These are currently areas of active research.

Having said that, we should right away reverse ourselves somewhat and state that this collision operator (by itself) actually works acceptably down to a plasma parameter in the low multiples of 10 (more explicitly, that is when the Coulomb logarithm is greater than 2) for some cases. Furthermore, with some modifications, we can (and have, for this thesis) run the collision package even where the number of particles per Debye sphere is below unity, at least without the behavior becoming pathological. More on this in a moment.

There is one further complication with actually implementing this collision algorithm into OSIRIS, the solution of which turns out to somewhat simplify and clarify the collision operator, making it worth mentioning here. The complication arises from the fact that in OSIRIS and other PIC codes the charge of individual macro-particles can vary, even for a single species with a constant charge to mass ratio, q/m . For example, one macro-particle may represent 10 real particles while another macro-particle with the same charge to mass ratio may represent 10,000 real particles, and hence have a simulated charge 1000 times greater. The ratio of the charge of the macro-particle to that of a real particle is termed the “weighting” of the macro-particle – though it is more often the relative weighting rather than the absolute weighting which is important (and in fact outside of collisions there often is no sense of absolute weighting). An example, where the use of variable weights for a single species would be useful is in a large density gradient. It is not desirable to have too few particles per cell in one region nor too many in another, so to change the plasma density we change the weighting on the particles in the different regions. Another example would be a plasma with ions at a high ionization state – although there are fewer ions than electrons, we might want to keep their simulated particle counts similar for better statistics and/or to reduce numerical noise and/or macro-particle collisions. Since the transformation to the center-of-momentum reference frame relies on the masses of the actual particles, if

we exchange momentum between the macro-particles (with different weights) in this frame neither energy nor momentum would be conserved. (Translating to the center-of-momentum frame for the macro-particles would be even worse, as that's not physically meaningful.)

Nanbu and Yonemura [57] proposed solving this problem by retaining the kinematics of the physical particle collisions for the macro-particles, but only updating the momentum of each partner in the pair differently based on statistics related to the particle weightings. More specifically, they construct a sum of the weights of the lowest weighted particle in each pair, the collision time-step is then adjusted in that cell using this sum. They then proceed with the collision calculation as before using this adjusted time-step, but when they update the momentum they do so randomly. Within each pair the particle with the smallest weight, W_s , is scattered, but the particle with the largest weight, W_l , is updated with probability W_s/W_l . In this way momentum and energy are not conserved in each collision event, but they are when averaged over many collisions.

An added benefit of this method of calculating the collision time step (though not obvious from this necessarily abbreviated description) is that it removes the need to use the arbitrary lesser of the two densities when calculating the collision frequency. This arises from the fact that in this formulation when colliding particles of different species, the species which defines how many pairs there are (and so defines the sum of weightings, call it species A) is always the species with more simulated particles, but the species which defines the density used in calculating the collision operator is the opposite species, species B. But using this method there is also an implicit calculation of the density of species A, which is the sum over the collision events. In this way species A collides off of species B, and species B collides off of species A symmetrically and there is no need to use or even ever calculate which species has the lower density. To use this method it is of course not necessary for the two species to have different weightings as the probability for scattering the largest weight particle in a pair is now unity, i.e., $W_s/W_l=1$. Therefore, this algorithm works seamlessly without any necessary modifications for colliding particles with equal or non-equal weights. It also actually addresses a problem we skipped over earlier: what to do when a species is colliding with itself and there are an odd number of particles. In this formulation, one particle, randomly picked from a

total of N , collides twice, and the collision time-step is reduced by a factor of $N/(N + 1)$. In OSIRIS we use the method of Nambu and Yonemura even for unequal weighted particles. Therefore, energy and momentum are not rigorously conserved in each collision but are only conserved in an average sense (when particle weightings are the same, energy and momentum are rigorously conserved).

Another – now physical – complication arises in extending these results to particles having relativistic speeds. That the results will be inaccurate for relativistic particles is readily apparent in the fact that no distinction has been made between velocity and momentum divided by mass, and evaluating the details of the interaction physics using relativistically correct calculations is non-trivial. However, modifying our Monte Carlo collision model to be consistent with relativistic physics turns out to be simpler (assuming Coulomb force interactions). Of course all frame transformations have to be Lorentz invariant; but otherwise the kinematics turns out to be unaffected, since the interaction is a rotation in momentum space, which remains valid. It therefore is only necessary to find what the modification is to collision frequency or strength of the collision operator (i.e. $\langle \delta^2 \rangle$ in the Takizuka and Abe model.) There have been a number of different methods proposed to address this ([58, 59, 60].) To our mind the most elegant approach is that of Pérez *et al.* [60]. They calculate the momentum-transfer cross section from the differential cross section, which they use as an effective total cross-section, σ , (defined in the center-of-momentum frame); then they exploit the fact that, for any total cross-section, the product $\sigma v_{rel} \gamma_1 \gamma_2$ is a relativistic invariant (where the γ 's are the relativistic boost factors for the two colliding species, and v_{rel} is the relative velocity). Using all this they arrive at a transformation of the effective collision rate, eq. 1.3, as seen in the lab frame, which can then be used in an otherwise unmodified Nambu calculation of the total deflection angle¹⁶. The interested reader is directed to consult Pérez *et al.* [60] for the details.

Let us now return to the somewhat opposite limit, that is a very cool plasma. In this case the model also breaks down, because the plasma parameter becomes small. There are

¹⁶ It is also necessary to modify the Coulomb logarithm for relativistic particles (in fact actually for particles where $\beta > \alpha$, where α is the fine structure constant) because the large angle scatterings can no longer be modeled classically.

two effects to take into account in this regime. First, the Coulomb logarithm shouldn't be allowed to become too small; it is customary to take the limit the lowest value to be 2. This is a semi-empirical result, but it can be qualitatively understood by observing that as the number of particles in a Debye sphere becomes very low, their ability to screen out charge becomes limited by more than their thermal motion, and so fields can extend farther than a Debye length before dying out. Second, because of the third power of the velocity in the denominator of the collision frequency, the collisionality diverges for very low temperatures¹⁷. More specifically, as the plasma parameter goes below unity, the mean free path of the particles becomes less than the inter-particle distance, which is clearly unphysical. Pérez *et al.* address this by clamping the collision frequency such that this isn't allowed to happen. The nature of the collision interaction stays the same but the strength goes to zero with decreasing temperature. Obviously this is an incomplete modeling of low-temperature physics, but it allows us to retain at least some sense of cold collisionality in our PIC code without incorporating new collision operators and importantly without the need to switch between models with if tests.

We close this section by commenting that it is our opinion that there is still work left to be done when combining the traditional PIC method with a Monte-Carlo method and to use PIC by itself to address questions regarding collisions in strongly turbulent plasmas. For example, the PIC method itself conserves energy very well (particularly when using higher order particle shapes). In addition, the Monte Carlo method by itself conserves energy well in an average sense. However, there is some evidence that combining these two methods together might lead to lack of energy conservation. The possible interactions between these methods is an area that deserves more attention [61]. As we noted earlier, it is possible to carry out the Lenard Balescu approach to study collisions including the effects of finite size particles. One could then run the PIC code and see how well the results are consistent with such a collision model particularly in the presence of strong turbulence. This could be aided

¹⁷ We note in passing that the nature of the divergence, or even if it diverges at all, can depend on the exact formulation of the Coulomb logarithm (if the argument of the logarithm is allowed to go to zero) – as all of the derivations of this term assume a large argument, they don't in general have consistent behavior when it becomes small.

by following the tracks of particles to watch their diffusion in momentum space. This is also a topic that deserves future study.

1.5 Resistive Collimation

A noteworthy aspect of ultra-high intensity laser-matter interactions – and an implicit assumption of the Fast Ignition scheme – is that they involve normalized and absolutely large currents of fast electrons. For FI, rough calculations using energy requirements and fuel disassembly times give current estimates in the range of hundreds of mega-amperes up to giga-amperes [62]. This is interesting because the Alfvén current – the highest current of electrons which can propagate in any medium [63, 64] – is approximately 50kA, orders of magnitude lower. At first this seems a contradiction, and perhaps FI is doomed from the start; however it turns out that, through a combination of charge separation and induction, a return current almost equal in magnitude, but of opposite sign, to the fast electron current is generated by the plasma, so that the *net* forward current satisfies the Alfvén limit [62]. Modeling the physics of such return currents is essential in order to get its global dependence correct and therefore it can be important to use isolated targets [65].

In FI we therefore have a total (net) current that may not be that high, though both the forward and return currents are themselves very large and much greater than that which is often encountered. In addition, in order for there to be a large net heat flux, the average kinetic energy of the electrons making up the forward and return current are widely separated, meaning that in some respects they can be treated as different species with qualitatively different physical behavior. In particular, the effect of Coulomb collisions can be quite strong on the thermal electrons drifting back towards the source, but almost absent for the relativistic forward “beam”. We use “beam” in quotes because the distribution function is a monotonically decreasing function of energy. The distribution function is asymmetric with the forward going electrons have higher energies than the backward flowing electrons. The combination of these two facts leads to a very interesting phenomenon that might arise in these interactions under certain conditions, in particular, when high densities and/or high

atomic number ions are present, which is known as *resistive collimation*.

To derive the equation for the growth of the magnetic field in a resistive medium, we will start with a generalized (but still simple form) Ohm's law from so called single-fluid plasma theory [66]. Often only the collimation term ([67, 68, 69, 70, 71, 72]) or the collimation and the diffusion term ([73, 74, 75]) are retained in this analysis. Even with a simple form for Ohm's law this leads to five terms, four of which we see as potentially relevant to FI electron beams (diffusion we will argue is not). However, we still find it necessary to make a number of further simplifying assumptions, in order to make interpretation meaningful. They are: 1) that the convective derivative of the total current is zero; 2) that the ratio of the ion to electron mass goes to infinity, $m_i/m_e \rightarrow \infty$; charge neutrality, so that $n_e = Zn_i$; 3) that the total pressure is given by the electrons, and is a scalar, p_e ; and finally 4) that the current can be broken into a component from fast electrons and one from the thermal background electrons, $\mathbf{J} = \mathbf{j}_{fast} + \mathbf{j}_{thermal}$, and that the collision term is well approximated by an effective collision frequency, ν , times the current of thermal electrons.

In addition, we will be using Ampère's law, in which we will assume that the displacement current can be neglected, $\partial \mathbf{E} / \partial t = 0$. Please not though that, although this approximation precludes space charge effects such as those that arise in streaming instabilities and precludes radiative effects, it does permit current filamentation and Weibel like instabilities.

With all this taken into account, the simple Ohm's law is

$$\mathbf{E} + \frac{\mathbf{v} \times \mathbf{B}}{c} = \frac{\mathbf{J} \times \mathbf{B}}{n_e e c} - \frac{\nabla p_e}{n_e e} + \eta \mathbf{j}_{thermal} \quad (1.6)$$

where $\eta = \nu m_e / n_e e^2$. Next Ampère's law without the displacement current can be written as,

$$\nabla \times \mathbf{B} = \frac{4\pi}{c} \mathbf{J} = \frac{4\pi}{c} (\mathbf{j}_{fast} + \mathbf{j}_{thermal});$$

which can then be used to solve for $\eta \mathbf{j}_{thermal}$

$$\eta \mathbf{j}_{thermal} = \frac{c}{4\pi} \eta \nabla \times \mathbf{B} - \eta \mathbf{j}_{fast}.$$

Last, we use Faraday's law, $\frac{\partial \mathbf{B}}{\partial t} = -c \nabla \times \mathbf{E}$, to find the equation for the evolution of the

magnetic field

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times \left(\mathbf{v} \times \mathbf{B} - \frac{c^2}{4\pi} \eta \nabla \times \mathbf{B} - \frac{\mathbf{J} \times \mathbf{B}}{n_e e} + c \frac{\nabla p_e}{n_e e} + c \eta \mathbf{j}_{fast} \right). \quad (1.7)$$

The first three terms include $\mathbf{B}$ itself, and therefore represent the modification or transport of an existing magnetic field. The first term transports the field with the bulk movement of the plasma, which can be seen by combining eq. 1.7 with all other terms on the right-hand side dropped, the continuity equation, $\partial \rho / \partial t + \nabla \cdot (\mathbf{v} \rho) = 0$, and the definition of the convective derivative, $D/Dt = \partial / \partial t + \mathbf{v} \cdot \nabla$. This leads to

$$\frac{D}{Dt} \frac{\mathbf{B}}{\rho} = \frac{\mathbf{B} \cdot \nabla}{\rho} \mathbf{v}.$$

In two dimensions this is even simpler, as for self-generated fields, $\mathbf{B} \cdot \nabla = 0$, and so the ratio of $\mathbf{B}$ to ρ can be seen to be a convective constant. The second term can be rewritten as $-\nabla \times (\frac{c^2}{4\pi} \eta \nabla \times \mathbf{B}) = \frac{c^2}{4\pi} (\eta \nabla^2 \mathbf{B} - \nabla \eta \times \nabla \times \mathbf{B})$, which, ignoring any gradient in the resistivity, gives a diffusion equation for the field; importantly this effect disappears when the resistivity goes to zero. The third term can be rearranged to give (again taking advantage of $\mathbf{B} \cdot \nabla = 0$ in 2D) to give $c \nabla n_e \times \nabla |B^2| / (4\pi n_e^2)$, and so can give an evolution of the field when a gradient in magnetic pressure and material density are transverse, and the magnetic pressure is sufficiently large.

The last two terms are the only terms which don't depend on $\mathbf{B}$, and so it is these terms that can create magnetic field when there was none before. The first can be expanded, using $p_e = n_e T_e$, to give $(c/e) \nabla n_e \times \nabla T_e / n_e$, and so can create magnetic field where temperature and density gradients are crossed. Such a situation arises, for instance, when a beam which is heating the background crosses a transverse material boundary. Note that an anisotropic pressure tensor can also create a B-field even if $\nabla n_e = 0$.

Finally the last term, which comes from the solenoidal part of the resistive electric field, can lead to the creation of a magnetic field if either the current or the conductivity have spatial dependence. A few points about this term are worth emphasizing at this time. First notice that for a finite width beam and a uniform collisionality, the sign of the field growth is the correct one to focus and channel the incoming fast electrons, and hence this term is

said to lead to magnetic collimation. However also notice that for the return current – the current which created the resistive field – the sign is opposite of that, and so would instead lead to a defocusing of the electrons in the return current. However, the return current (with its lower velocity particles) is not strongly effected by the magnetic field. This is why having the two distinct currents, with distinct collisionalities, is a unique and key feature of resistive collimation. Lastly, there are actually two effects in the last term as can be seen by rewriting it as: $\nabla \times (c\eta \mathbf{j}_{fast}) = c\eta \nabla \times \mathbf{j}_{fast} - c\mathbf{j}_{fast} \times \nabla\eta$; the first term will lead to collimation of a cylindrical beam as just stated; however, taking into account the Ohmic heating of the background and assuming Spitzer resistivity (in point of fact due to the return current, though the result is identical either way), the second will oppose the first and lead to a decrease in collimation or even beam deflection. The non-linear interplay between these two effects, even ignoring the other terms in the Ohm’s law, let alone the actual kinetic nature of the collision term, shows that resistive collimation is a very complex phenomenon, for which analytical models struggle and even in simple situations numerical modeling is key.

1.6 Outline of Disseration

The rest of the dissertation is broken into four chapters.

Chapter 2 analyses and discusses the basic physical mechanisms responsible for laser absorption at a sharp interface between vacuum and a highly overdense plasma. Such physics is important for converting intense and energetic lasers into relativistic electrons as is needed in the FI concept. In particular, by using one- and two-dimensional PIC simulations we will show that there are two main mechanisms present in these interactions, one which happens directly at and in front of the material surface in a single laser cycle and another which happens chaotically over a variable length of time throughout an underdense region of plasma that forms in front of the target over ion time scales. A key observation developed in this chapter is that it is the former mechanism which generates the moderate energy electrons necessary for effective coupling to core in the FI scheme, while the latter leads to electrons of unbounded energy and hence is inefficient from an FI perspective. Thoughts on how to

optimize the surface acceleration and minimize the underdense acceleration are presented.

Chapter 3 discusses an experiment on the Omega EP laser system and PIC simulation modeling of that experiment. The design on the experiment was to look for differences in the transport of an electron beam – generated using an ultra-high intensity laser focused onto a solid (gold) target layer – through either cold matter or a pre-shocked (warm) material in what is sometimes known as the ‘warm-dense matter’ state. The experimental results were that total beam transport – as measured in particular through K- α radiation from a copper foil on the far side of the target – was greatly diminished in the warm-dense matter case as compared to the cold case. Our PIC modeling indicates that a prime cause of this different behavior is likely to be the difference in the magnetic field structure created in the two types of targets. In the case of cold matter a quasi-static magnetic field is formed which leads to resistive collimation of the forward beam. Contrariwise, in the warm – or as it turns out more importantly, less dense plasma – no such collimating field arises. And, in addition, the B-field which *is* created, via the current filamentation instability, is transported back with the material current to edge of the gold layer, creating a strong magnetic field at the gold/plastic interface field which potentially disrupts the forward current.

Chapter 4 presents theoretical analysis and simulation of a numerical effect which we term ‘Macro-Particle Stopping.’ This effect is caused by the ‘macro’ nature of the particles in a particle-in-cell code, where one macro simulation particle can be thought of as representing many physical particles. In particular, we find that the effect of the electrostatic wake of a relativistically traveling particle is proportional to q^2/m , rather than to q/m as would be the case for the response to an externally created field (that ratio being independent of macro-particle size). This leads to anomalous stopping of the particle, and simultaneously an anomalous heating of the background electrons. We find that this effect is very relevant to the parameters often used in FI simulations, both in our own work and in many other published results. We also present and demonstrate a solution for reducing the effect of this macro-particle stopping to acceptable levels, by splitting the highest energy particles (a technique which was used in almost all the simulations in chapters 2 & 3). We reran our global integrated FI simulations previously published Tonge et al, uses the splitting algorithm

to better understand the fast electron spectra for FI.

Finally, **chapter 5** summarizes these results, provides conclusions, and discusses applications and directions for future work.

CHAPTER 2

High Intensity Laser Absorption in One- and Two-Dimensions

The mechanism of absorption of intense laser light by overdense matter is of interest to a variety of experimental and technological applications,

Fast Ignition Inertial Confinement Fusion and Compact Ion Accelerators. In addition it is interesting from a purely scientific perspective, as a collisionless plasma phenomena in a relativistic regime. As more and higher quality high intensity laser facilities come online, it will be of increasing importance to have an understanding of the underlying physics of high-intensity laser-matter interactions, both to plan experiments correctly, and to be able to match the diagnostics from experiments to the underlying physical processes which generate them.

Various mechanisms have been proposed for the absorption of laser light in this regime [47, 76, 46]; some mechanisms may take place simultaneously, while others are mutually exclusive. Three of the key factors which can affect the absorption of laser light in high density plasmas are, the distribution function of the plasma at the surface, especially the transverse momentum; the spatial distribution of underdense plasma in the region in front of the critical density surface (the so-called ‘pre-plasma’); and the transverse structure of the critical surface, especially at scale-lengths comparable to the laser wavelength.

In this chapter, we examine the absorption of laser light in a regime where the pre-plasma is minimal, where the surface of the plasma is initially flat, and where the transition from underdense (in fact vacuum) to greatly overdense is over a scale much less than the laser wavelength. This situation is very relevant to high intensity lasers with minimal prepulses

interacting with overdense plasmas. We have discovered that the standing-wave structure of the superposition of the incident and reflected light is critical to the mechanism of acceleration of electrons, with both the electric and magnetic fields playing key roles.

Although the absorption of high-intensity laser light potentially has many applications, probably even including some which haven't been imagined yet, the initial and primary motivation of this research was in the context of the Fast Ignition concept. As such, most effort was put in to discovering and optimizing the mechanism which generates high fluences of moderate energy electrons, which are believed to be most able to couple to the compressed core of a fusion target – usually estimated to be those in the range of 1-3 MeV, that is with a relativistic γ factor in the low integer range. This is the reason for the emphasis on the regime of sharp, flat surfaces: as we will show this is the ideal - in fact essential - geometry for coupling the laser efficiently to the plasma without generating very energetic, "hot", electrons of unbounded energy.

2.1 1D Simulations

We begin by looking at simulations in only one spatial dimension. We do this because it is the most simple geometry and hence the easiest to analyze; it also clearly illustrates the basic features of the interaction mechanism which can be harder to isolate in the more complicated and multi-faceted regime of multi-dimensionality.

However we would also like to emphasize that some of the results, even some of the most essential, will be erroneous and contradicted by what is found in higher dimensions. Specifically the strong dependence on initial temperature is not a feature seen in 2D or 3D simulations (with the caveat that in 2D *s-polarized* simulations the results do depend strongly on the initial temperature). The reason for this is that in 2D p-polarized and full 3D simulations the plasma heats up strongly for unavoidable reasons that will become clear soon. In these cases the results do not therefore depend on the initial temperature. A number of groups have attempted to explore the absorption of intense laser light in one dimension, both in simulation and theory [46, 47], most likely because it is most computationally efficient

to do so. Their results, though, are in general either not present in higher dimensions or are at best of transient and marginal importance to the absorption mechanism. The reason is that – as we will show briefly – although parts of the absorption process have one (spatial) dimensional models, the actual response of the plasma is the result of the constant interplay and equilibrium of single- and multi-dimensional physics. Therefore the results that follow presently should be thought of as investigating the individual parts of the interaction singly, and decidedly as *not* modeling an actual laboratory situation, even an idealized one¹.

2.1.1 Fixed Ions

A typical simulation setup in one-dimension is as follows. The plasma is initialized inside a box $75.6\ c/\omega_0$ long, where ω_0 is the frequency of the laser; this is approximately 12 laser wavelengths long, which of course is $12\mu m$ assuming a $1\mu m$ laser. To the left of the plasma is a $6.3\ c/\omega_0$ ($1\mu m$) vacuum region. At $x_1 = 6.3\ c/\omega_0$ the density of the plasma ramps up as a step function to $100n_c$ ($\sim 1.11 \times 10^{23} cm^{-3}$ for a $1\mu m$ laser, or $\sim 1.006 \times 10^{23} cm^{-3}$ for a $1.053\mu m$ laser) and remains constant for the rest of the simulation box. At time $t = 0$ a laser pulse is launched from the left-hand side of the box, with a normalized vector potential a_0 of 6 (which gives an intensity $I \sim 4.93 \times 10^{19} W cm^{-2}$ for a $1\mu m$ laser), and is polarized with the electric field in the x_2 direction. The rise time is $12.6/\omega_0$ (two laser periods, $\sim 6.67 fs$), after which it stays constant. Initially we consider a perfectly cold plasma, i.e. $v_{th} = 0$ ². This is admittedly unphysical, but as we will see shortly temperature – and especially, transverse temperature – plays a key role in the laser absorption, and so it is best to remove the effects of the initial temperature when performing the initial analysis.

The numerical parameters are as follows. The cell size is $0.0105c/\omega_0$ (sufficient to adequately resolve any skin-depth scale physics), there are 600 particles per cell, and the time step is set by the Courant condition. The particle shapes for both current deposit and parti-

¹ In fact the one-dimensional simulation we show which is probably the most pertinent to the ‘real’ result is the one with an initial plasma temperature of 511 keV, and the authors don’t think they need to work hard to be convincing that such a target (at solid density!) is not a very reasonable one by itself.

² In most PIC codes it is difficult to look at low temperature phenomena due to spurious heating caused by the grid (‘grid heating’). However in OSIRIS we are able to greatly reduce this effect by using higher order particle shapes.

cle push are third-order splines (‘cubic’); before advancing the fields the current is smoothed with a 4-pass (1,2,1) filter followed by a (-5,14,-5) compensator; and the field solver uses a standard Yee [77] mesh. Boundary conditions are absorbing (Lindman [78]) for fields, and specular reflection for particles. A charge conserving current deposit is used [79]. Also, as an ion species is not initialized and the initial electric field is set to zero, OSIRIS behaves as if there is an infinitely massive neutralizing ion background.

The phase space of this simulation after ~ 250 fs of laser illumination³ is shown in figure 2.1. Figure 2.1.a shows a 2D colormap of particle count versus forward position (along the laser propagation direction) and forward momentum. One can see a slightly warmed, isotropic plasma around $p_1 = 0$;⁴ a thermally broad but rather spatially uniform forward-going group of electrons with a few orders of magnitude less density, with velocities up to about $0.15c$; and an even sparser beam tail above that with velocities peaking at about $0.3c$, with perhaps a small amount of spatial coherence. One can also see that no particles have escaped into the vacuum region in front of the target, even at this relatively late time. Figure 2.1.b shows the distribution function, found by spatially averaging a small section of 2.1.a (in the range $8 < x_1 < 10$); it shows the same groups of electrons: a thermal background of about $v_{th} = 0.01c$, and a beam with momentum out to $p_1 = 0.3m_e c$, with a break in slope between them at $p_1 = 0.15m_e c$. Thus we find that no particles are accelerated to more than 5% of the quiver velocity of the laser (which is given by $p = a_0 m_e c$, i.e. is 6), and most are confined to 2.5% or less. Figure 2.1.c shows the total heat flux carried by the electrons as a function of $\gamma - 1$ (kinetic energy) – averaged in the same volume – which shows that the bulk of the energy is carried by particles with a $\gamma - 1 < 0.01$ (corresponding to a momentum $|\vec{p}|/m_e c < 0.1$), even though such particles both carry less energy each and have lower velocity. Calculating the energy absorbed by comparing the forward heat flux of the

³ Most of the the one dimensional simulations with similar box sizes will be run for a shorter period of time, to avoid allowing particles to reflux off the back of the box and interacting with the laser multiple times. (Alternately, as we will see in the two-dimensional simulations, the refluxing can be eliminated with careful use of more advanced boundary conditions; c.f. appendix B). However in this case the velocities are non-relativistic, so acceptable timescales are longer.

⁴ Careful inspection of 2.1.a or 2.1.b reveals that the cold plasma is not in fact centered around $p_1 = 0$ but actually around a slightly negative p_1 . This is due to the bulk plasma setting up a return current to counteract the current of the hot electrons and maintain quasi-current neutrality for the system.




FIGURE 2.1: Phase space data for a 1D simulation of an $a_0 = 6$ laser incident on a plasma with the initial temperature of $T_e = 0$. Sub-figure 2.1.a shows the $p_1 x_1$ phase space (i.e. particle count vs. forward position and forward momentum); sub-figure 2.1.b shows the forward momentum distribution function at the front of the plasma, that is essentially a line-out of figure 2.1.a (though actually an average over a small section); sub-figure 2.1.c shows the forward heat flux as a function of the kinetic energy ($\gamma - 1$) of the electrons carrying it. A small excess of forward-going electrons can be seen in 2.1.a and 2.1.b, and some net heat transport in 2.1.c; but the peak momentum is $\sim 0.3m_e c$, i.e. the (proper) velocity is 5% of the quiver velocity, $u_q = a_0 c$, and total heat flux is low and carried by the lowest energy particles.

electrons in a spatially ($8 < x_1 < 10$) and temporally (one laser period) averaged window inside the plasma⁵, we find that total laser absorption is 0.0145%. Clearly absorption and energy conversion are very low for a cold plasma, at least in this geometry.

We next look at a slightly more realistic simulation, where all parameters remain the same except we give the electrons a finite temperature. We will initialize the plasma as a Maxwellian with a proper velocity of $1.0 \times 10^{-1}c$ in the x_1 and x_2 directions, and zero in the x_3 direction⁶; this is equivalent to a temperature of approximately $5keV$. Figure 2.2 shows the same phase space diagnostics as are shown in 2.1 after $\sim 50fs$ of laser irradiation. Looking at the 2D p_1x_1 phase space in 2.2.a, the most noticeable feature is the existence of bunches of high γ particles streaming into the plasma (appearing as upward pointing ‘strings’, ‘arms’, or ‘fingers’ in this diagnostic). These contain electrons with momenta up to almost $3m_ec$, i.e. one half of the quiver velocity. They are spaced about $\pi c/\omega_0$ apart (half a laser wavelength), indicating that they are generated once every half laser cycle because they have a speed near the speed of light. They also have the characteristic ‘leaning into the wind’ appearance of high energy particles out running lower energy ones, indicating that they all initially enter the plasma together (a conclusion reinforced by the vertical appearance of the very first bunch). However the arms are very thin and tenuous, a point made more clear by looking at the spatially averaged distribution shown in figure 2.2.b (again averaged in the range $x_1 = [8, 10]$, i.e., centered around the second ‘arm’) which indicates that the total density contained in a beam is orders of magnitude less than the background plasma density. Also note that once again exactly zero particles have been able to escape into the vacuum region.

⁵ Although it would be more direct to measure absorption by looking at the Poynting flux, the net signal can be very weak compared to peak value; in addition there is a beat between the incident and reflected beams, and there are harmonics in the reflected beam. The heat flux does not have these complications, and so is usually the preferred diagnostic. (The added complication is that any absorption which leads to heating at the front of the target would be missed; however these measurements are taken after a pseudo-steady state is reached at the surface, so heating should be minimal.)

⁶ It can easily be shown that due to the conservation of transverse canonical momentum in 1D geometry (see the appendix), the physical momentum in the direction perpendicular to the laser polarization is strictly conserved at all times. Although setting a temperature in this direction would have virtually no effect on the physics either way (with the interesting caveat that for relativistic temperatures, it would appear to modify the rest mass of the electrons). we have a slight habit of making it zero in such cases, as it somewhat reduces the noise.




FIGURE 2.2: Phase space data for a 1D simulation of an $a_0 = 6$ laser incident on a plasma with the initial temperature of $T_e = 5keV$. Sub-figure 2.2.a shows the $p_1 x_1$ phase space; sub-figure 2.2.b shows the forward momentum distribution function at the front of the plasma; sub-figure 2.2.c shows the forward heat flux as a function of the kinetic energy of the electrons carrying it. Bunches of accelerated electrons can be seen in 2.2.a, with the peak momentum of roughly $3m_e c$, that is half of the quiver velocity $u_q = a_0 c$. The bunches however are very low density, and total heat flux remains small; as in 2.1.c the heat is carried predominantly by low energy electrons.

As the initial plasma has a much broader distribution than that seen in 2.1.a, it's harder to see the moderate energy ($p_1 \sim +0.3m_e c$) beam just above the thermal energy; none the less if you look closely at 2.2.a it's apparent that the bulk distribution is not symmetric about the momentum peak and is broader in the forward-going direction. This conclusion is confirmed by looking at the heat flux versus γ in 2.2.c, where it is clear that most of the energy is being carried by particles with kinetic energies around $0.1m_e c^2$, and very little is carried by the high energy electrons. Measuring the spatially and temporally averaged energy flux being carried by the plasma, we find that total energy absorption is 1.11% of the incident laser energy. Therefore we find that although the energy absorption is about 2 orders of magnitude greater in a 'realistic' one dimensional plasma than in a cold plasma – and although we are able to accelerate some electrons to relativistic energies – total laser absorption still remains quite low, and what little energy is absorbed is carried almost entirely by non-relativistic electrons.

Next let's look at another somewhat unphysical simulation but in the opposite sense to that shown in figure 2.1, i.e. the plasma has an unreasonably high temperature. Again we keep all other parameters the same, but initialize a plasma with a thermal proper velocity of $1.0c$ in the x_1 and x_2 directions, which gives a temperature of roughly $313keV$. Results are summarized in figure 2.3, showing the same diagnostics as figures 2.1 and 2.2. Looking at figure 2.3.a, we see features similar to figure 2.2.a, but with notable differences. Most apparent is that there are now electrons all the way up to twice the quiver 'velocity,' with proper velocity $u_1 = 12c$. Also the 'arms' are now bent back towards the surface rather than leaning into the plasma, appearing almost like 'hooks' from a microscopic picture of Velcro. This shape, coupled with the highly relativistic momenta (i.e. all velocities are near c), indicates that the highest energy electrons enter the plasma at a later time than the lower energy front of the bunch. The distance between the front of the bunch and the highest energy electrons is just about $\pi c/2\omega_0$, indicating that a quarter laser period separates their entrance into the plasma. The electrons which enter last make the downward arch of the hook. A closer examination of each hook reveals that the density of each gradually decreases as one moves from its base up to the tip of the hook. This indicates that fewer electrons




FIGURE 2.3: Phase space data for a 1D simulation of an $a_0 = 6$ laser incident on a plasma with the initial temperature of $T_e = 316 \text{ keV}$. Sub-figure 2.3.a shows the $p_1 x_1$ phase space; sub-figure 2.3.b shows the forward momentum distribution function at the front of the plasma; sub-figure 2.3.c shows the forward heat flux as a function of the kinetic energy of the electrons carrying it. Bunches of electrons with velocities up to twice the quiver velocity can be seen going into the plasma. The density of these bunches is nearly constant from $p_1 = 5 m_e c$ up to the peak momentum, as can be seen in fig. 2.3.b. Similarly the heat flux is partitioned close to equally among the energetic particles, as indicated by fig. 2.3.c being almost flat between $(\gamma - 1)$ of 2 and 12. Circled particles and areas indicated by arrows are explained in the text.

re-enter the plasma at the end of each cycle. One can also see that, in contrast to lower energies, a small number of electrons can escape into the vacuum region (they appear here as very faint dots, a select few are indicated by circles), and there are also now high energy electrons in the plasma which are not part of the hooks. It will turn out that these two facts are related. Furthermore, although also difficult to discern on this color map, there are small ‘divots’ in the bulk plasma at the base of each hook (marked by arrows), indicating fewer electrons compared to the adjacent phase space on either side.

Although still quite narrow in x_1 at a given energy, the arms do seem more substantial (wider, higher, and denser) in figure 2.3.a than in figure 2.2.a. This can be seen more quantitatively in figure 2.3.b, where again the distribution is averaged over the window $x_1 = [8, 10]$, i.e. around the second arm. It is apparent that although still orders of magnitude less dense than the plasma, the very energetic electrons, all the way up to $p_1 = 12m_ec$, are now relatively large in number, indicating that the relativistic electrons now play a key role in the dynamics. This is further confirmed by looking at the forward heat flux in the same region in figure 2.3.c, which shows that the energy is now carried equally by relativistic electrons all the way out to the maximum energy. Although the momentum distribution is now far too broad to easily see any shifts in the bulk plasma, this effect can be seen instead also in figure 2.3.c as a strong negative peak in heat flux for the lowest energy electrons. Measuring the spatially and temporally averaged total heat flux, we now find that the energy carried by the electrons is equal to 24.8% of the incident laser energy; the time-averaged Poynting flux gives the same result to within a few percent. We can therefore conclude that for a sufficiently hot plasma, there can be significant energy absorption and significant relativistic electron generation by a high a_0 laser at a steep interface, even in one dimension.

We have shown through sample simulations that the initial temperature of the plasma plays a key role in the laser absorption at steep interfaces, with the absorption varying by orders of magnitude between the cold ($u_{th} = 0$), ‘tepid’ ($u_{th} = 0.1c$), and hot ($u_{th} = 1.0c$) results in one dimension (though none of these individually represent an adequate modeling of a correct, multidimensional result). In addition, as we stated earlier it is specifically the *transverse* temperature which affects the absorption. This is especially important be-




FIGURE 2.4: Phase space data for a 1D simulation of an $a_0 = 6$ laser incident on a plasma with the initial temperature of $T_e = 5\text{keV}$ in the forward direction only, and $T_e = 0$ in the transverse direction. Sub-figure 2.4.a shows the $p_1 x_1$ phase space; sub-figure 2.4.b shows the forward momentum distribution function at the front of the plasma; sub-figure 2.4.c shows the forward heat flux as a function of the kinetic energy of the electrons carrying it. Comparing to fig. 2.2 (which has the same temperature but isotropic), the accelerated electrons are clearly greatly diminished in both density and peak momentum; the ‘bunches’ of electrons can barely be called that any more, as they stream over each other and dissipate even in the short distance shown. Correspondingly the total heat flux is significantly lowered.

cause in one-dimension transverse canonical momentum conservation means that it is in fact impossible to heat the plasma in the transverse direction, so that if it is this temperature that controls the absorption mechanism, then the total absorption will depend on an input parameter which is chosen more or less arbitrarily at the start of the 1D simulation.

Unfortunately it isn't possible to have a plasma warm only in the ignorable direction as it is quickly and strongly unstable to the Weibel instability [80]. However it is of course perfectly stable to have the opposite, that is to have finite temperature in the physical direction and zero temperature in the ignorable direction; we thus take advantage of this to isolate the physics and carry out a simulation to see the effects of longitudinal temperature alone. The results are shown in figure 2.4. We use $u_{th} = 0.1c$. Figure 2.4.a is shown with the same parameters as figure 2.2.a to allow easy comparison. The 'strings' spaced at half the laser wavelength are clearly greatly diminished when the transverse temperature is zero, both in density and maximum momentum; in addition the moderate energy group of electrons just above the thermal background at $p_1 \sim 0.3m_e c$ in figure 2.2.a is not discernible at all in figure 2.4.a. Comparing the distribution in figure 2.4.b with that in 2.2.b one sees similar differences, with a sparse beam just barely above an almost unperturbed thermal background, and figure 2.4.c shows a barely measurable amount of net heat flux at modest γ . Calculating the laser absorption from the total heat flux gives a net absorption of 0.0735%, a factor of 5 greater absorption than the isotropically cold case, but a factor of 15 less than isotropically warm. Since the difference in total kinetic energy between isotropically and anisotropically warm is only a factor of 2, it is therefore clear that the transverse momentum has the largest effect on the laser absorption. The initial longitudinal temperature must necessarily be present to keep the plasma stable against the Weibel instability, but it is not needed for absorption.

To obtain a more functional understanding of the dependence of laser absorption on plasma temperature, we carried out a set of simulations where the initial temperature was varied in steps on a log scale. Total absorption was measured via the same method as described above: by running the simulation until any transient behavior from the initial laser interaction died away (about $50fs$), and then calculating the temporally and spatially averaged heat flux being carried by the electrons into the plasma. Figure 2.5 presents the




FIGURE 2.5: *Total laser energy absorbed for an $a_0 = 6$ laser in one dimension, for different initial temperatures with either isotropically warm plasmas or plasmas with temperature only in the physical direction, shown as a log-log plot. Only the very coolest plasmas give absorptions which are not a function of the initial temperature, and it isn't till the temperatures become relativistic that absorption approaches unity. Anisotropic plasmas show significantly lower absorption than isotropic plasmas.*

results as a log-log plot for a set of 30 simulations. There are two simulation curves which correspond to $v_{th1} = v_{th2}$, ‘isotropic’ (red crosses); or $v_{th2} = 0$, ‘anisotropic’ (blue x’s). The $v_{th1} = v_{th2} = 0$ case is not shown, but it had the same absorption as the lowest points on the graph.

A number of results are apparent. First, for initial isotropic temperatures below $T_e = 5 \times 10^{-7} m_e c^2$, i.e. $2,900^\circ K$, or 0.26 eV , the absorption is independent of temperature and the result is similar to the idealized case of $T_e = 0$. This is in contrast to the work of Sanz et al.[47], who propose a cold-fluid theory to explain the absorption they see in a 1D PIC simulation with $T_e = 500 \text{ eV}$. Next, for temperatures between $T_e \sim 10^{-4} m_e c^2$ and $T_e \sim m_e c^2$, the absorption is given by a gentle power law; measuring the slope of the graph gives an exponent of roughly 0.72. As the temperatures become relativistic the absorption levels off as energy conversion approaches unity, approaching as high as 73.8% for $T_e = 10 m_e c^2$.

An examination of the anisotropic results shows that as has been mentioned previously and as we have seen in one case already, the absorption for this situation is much less than

the isotropic one. In general, when v_{th1} is the same, it is decreased roughly an order of magnitude; and so the temperature must increase multiple orders of magnitude to achieve the same absorption. It is also worth noting that the ‘turn on’ temperature for the anisotropic case, that is the temperature at which absorption starts to increase from the cold result, is equivalent to a thermal momentum of $p_1 \sim 0.01m_e c$. This is roughly the temperature we saw in the background plasma group in figure 2.1, indicating that the absorption in the cold case is in fact due to the small amount of *longitudinal* heating that is possible for that case.

2.1.2 Mobile Ions

At first glance it may seem that we have artificially restricted the physics when we implicitly stipulated that the ions be fixed by not initializing a neutralized species, creating an artificial model to show that the one dimensional result isn’t pertinent to the multi-dimensional situation. Furthermore, we seemingly did so rather unnecessarily in that it would be a mere doubling of computational cost to include an additional species⁷. One might easily imagine that mobile ions might provide a lower density ‘shelf’ in front of the main plasma. This would lead to a much higher absorption in the sub-critical density plasma, which would then transport heat into the over-dense plasma, thereby leading to a higher temperature and thus higher absorption. Again it is still only possible to heat in the longitudinal direction, meaning that at best mobile ions will give a slightly ‘less unphysical’ simulation geometry. . An argument along these lines was made by Kemp et al. in [46], although their logic and approach was slightly different, postulating that an ion shelf of specific density and depth would setup an electrostatic field which would then ‘pull’ electrons into the laser field. Temperature did not play a part in their explanation.

Indeed when we look at the one-dimensional cold plasma result with mobile ions, we do see somewhat of an ion shelf, greater electron heating, and indications of much greater laser absorption over the same simulation with fixed ions. However, as we will now show, this is

⁷ In actuality there is nothing in PIC linking the total number of macro-electrons and that of macro-ions, and we can and do choose the ratio freely; however *more* ions than electrons is rarely done in practice, making this for practical purposes an upper bound.




FIGURE 2.6: *Forward momentum-forward position phase space plots for electrons in a 1D simulation of an $a_0 = 6$ laser incident on a cold plasma with mobile ions also included (which have a mass-to-charge ratio approximately that of D^+ ions). Sub-figure 2.6.a shows the phase space in the entire simulation, 2.6.b shows a zoomed-in image of the surface region. Deep in the box some electrons up to $p_1 = m_e c$ can be seen, but no such particles are present near the surface. Although the plasma has been heated, especially at the front, except for this bunch there is hardly any asymmetry between the forward- and reverse-going parts of the distribution function.*

in fact primarily due to an entirely different mechanism than that present in the fixed ion case, and is no more relevant to the ‘correct’ result than other 1D simulations⁸.

To study one dimensional mobile ion laser absorption, we use a simulation setup identical to that shown in figure 2.1, except for two differences. We now include a mobile ion species (we use the same number of ions as electrons), with a mass-to-charge ration 4000 times times that of the electron, this being roughly the correct ration for deuterium⁹. The ions are also cold We also make the box much longer, as we need to run for longer to let the transient behavior die out and unlike the fixed ion cold case it is no longer true that there are no

⁸ We will see later on that it does bare some similarities to the multi-dimensional *circularly* polarized case, however these similarities are more intellectually interesting than they are actually instructive.

⁹ Interestingly, this is roughly the correct charge-to-mass ratio for virtually all fully-ionized nuclei besides hydrogen and tritium, and hence so long as collisions and other single particle effects can be ignored these results apply to any material.

relativistic electrons that can bounce off the back of the box and return to the front.

Figure 2.6 shows the phase space results for electrons. As can be seen in 2.6a, there are electrons traveling deep into the plasma with momenta up to and exceeding $p_1 = m_e c$. However, as there are not electrons nearer the surface of the plasma with such large momenta, it is clear that their generation was a transient event and is no longer occurring. In fact their generation was related to the evolution of an ion shelf – they are the electrons which were on that shelf. However once they are driven off by the laser the shelf remains bare and so no further electrons can be accelerated off of it. Looking at figure 2.6b, which shows a zoomed in region of phase space near the surface, one sees that the number of backward What is more striking is that the plasma has heated up considerably, especially in a roughly skin-depth thick layer right at the surface. Also noteworthy is that the plasma has compressed quite a bit – originally the electrons went to $x_1 = 2\pi$, the left side of the box in figure 2.6b.

One might therefore expect non-negligible laser absorption, especially if the temperature in the surface layer is the pertinent factor (which is measured to be roughly 50 keV¹⁰). Indeed, if one measures the absorption by calculating the spatially and temporally averaged Poynting flux in the vacuum region, one finds a value of 2.70% – still a small amount, but much more than the 0.0145% seen in the fixed ion case. And in fact that absorption level is rather consistent with an anisotropic temperature of 50 keV from figure 2.5. However as we discuss next that turns out to be a complete coincidence.

That something entirely different is going on is made clear by looking at figure 2.7, which shows the energy flow in the simulation in the material species. The magenta line shows the heat flux in the electrons, which has a significant value right at the surface of the plasma but then quite quickly dies down by an order of magnitude¹¹. Contrasting this is the heat flux carried by the ions, displayed in green, which shows significant energy being carried into the plasma. Looking at the ion momentum phase space, shown in figure 2.8, makes it clear what is actually happening: a compressional shock is being driven into the plasma, accelerating

¹⁰ Interestingly though the distribution in that region is no where near a Maxwellian but is in fact an almost perfect waterbag, as in addition is the slightly cooler region right behind it.

¹¹ And this spike does not in fact represent any absorption; rather it is the hot plasma - which was warmed in a transitory event - traveling inward as the plasma is compressed.




FIGURE 2.7: Material heat flux in the plasma. *Magenta: Heat flux in electrons. Green: Heat flux in ions.* The absorption physics is dominated by the contribution of the ions, as can be inferred from the fact that they carry most of the energy. In fact even the narrow spike in the electron heat flux near the plasma edge is not related to (current time) absorption, but rather is caused by the inward motion of the already hot surface.

ions to twice the implosion velocity as they bounce off the shock. Equating the momentum carried by these ions with the momentum imparted by the laser, it can be readily derived that the compression velocity of the plasma is given by [38]

$$\beta_{front} = \frac{a_0}{\sqrt{2}} \sqrt{\frac{n_0}{n_p} \frac{m_e}{m_e + M_i/Z} \frac{1+R}{2}}$$

or, assuming the ion mass, M_i , is large compared to the electron mass, and that reflectivity R is unity¹²,

$$\beta_{front} = \frac{a_0}{\sqrt{2}} \sqrt{\frac{n_0}{n_p} \frac{m_e Z}{M_i}} \quad (2.1)$$

For the parameters of the simulation this gives $\beta_{front} = 0.0067$, compared to a value of 0.0069 found by measuring the actual displacement of the plasma front. From eq. 2.1 the total power being absorbed by the ions can also be easily found, and is given by

$$Power/Area = \sqrt{2} a_0^3 \sqrt{\frac{n_0}{n_p} \frac{m_e Z}{M_i}} n_0 m_e c^3$$

¹² It is in fact possible to retain the reflectivity throughout the analysis and arrive at an analytic expression for it, though the final form is rather complex.




FIGURE 2.8: *Forward momentum-forward position phase space plots for ions in a 1D simulation of an $a_0 = 6$ laser incident on a cold plasma. One group of ions can be seen traveling inward as a bunch with the compression velocity of the surface ($0.007c$), and another group, which reflected off the inward-moving surface, can be seen traveling into the plasma as a beam with twice this velocity.*

or, normalized to the laser intensity

$$Absorption = 2\sqrt{2}a_0\sqrt{\frac{n_0}{n_p}\frac{m_e Z}{M_i}} \quad (2.2)$$

Using the values from the simulation, we arrive at a predicted absorption of 2.68%, compared to 2.70% as measured from the Poynting flux in the simulation.

Therefore we find that the enhanced absorption we observe in one-dimension when we allow ion motion is not due to the effects of electron heating, an ion shelf or other modification of the plasma surface, or any other complex mechanism which would allow for electron acceleration; rather, it is simply the mechanical consequence of momentum conservation when that momentum must be carried by a species with finite mass. Hence, a one-dimensional geometry remains a poor choice for studying electron acceleration, even with the addition of a realistic ion species.

2.2 Early Time 2D Simulations - Absorption at a Flat Surface

We have seen that in the one dimensional case, where the evolution of the plasma distribution and the laser-plasma interaction region are greatly limited by the constraints of the geometry, the absorption of the laser is primarily a function of the initial parameters of the simulation – electron temperature, ion mass, vacuum-overdense transition shape, etc – and depends only transiently on time before assuming a steady-state result. In fact the constraints are so great and consequently the results so strongly dependent on the essentially arbitrary choices of the simulationist, that it is our opinion that one dimensional simulations are of very little use in understanding real results. Attempts to either develop theory to explain one dimensional simulation results, or to develop one dimensional simulations to explain higher-dimensional simulation or experimental results, has led to a great amount of confusion and contradictory answers for the relevant absorption mechanisms.

The self-consistent evolution of the laser-plasma interaction in higher dimensions is quite different, with the plasma having a natural response which is only transiently dependent on initial conditions. We focus on the results in two-dimensional planar (‘slab’) geometry, which is naturally more computationally efficient and easier to analyze than three-dimensional results. Although of course the same can be said of one-dimensional simulations – and we are arguing strongly that that is not a good reason for using them – we still believe 2D simulations are sufficient in most ways and are a good compromise. In particular we will show that the ability of the plasma to heat transversely – forbidden in 1D but possible in 2 – is a key facet of the interaction. When we believe the planar-geometry may be insufficient and analysis in high geometry is warranted, we will say so in the text.

We first show results from a short simulation ($<100\text{fs}$) with an initially cold plasma and mobile ions. This is to demonstrate clearly that electron acceleration is possible, and that it happens early in time so that the ions have moved very little (results with fixed ions are up to this point in most ways similar, though the peak acceleration is less, for reasons which we will return to). This will also show that the phase space diagnostics of the accelerated electrons most resemble the one-dimensional result in the case of an initially hot plasma,

as stated earlier; and also that although the whole plasma is not as hot as it was in that case, the temperature at the surface is in fact comparable. With these facts in mind, we then return to a somewhat simpler simulation, with fixed ions and preheated electrons. This second part can be understood from 1D arguments and simulations. We use particle tracking in this simulation to explore the nature of the acceleration mechanism.

Based on these results, we propose a model to explain the results, which involves two parts: the initial heating of the plasma surface due to instabilities in the current profile, and the interaction of this heated layer with the standing wave structure of the reflected laser light. We then use a numerical calculation to demonstrate the applicability and robustness of the standing-wave model; and we look briefly at the heating mechanism in theory and using a modified version of OSIRIS. Lastly in this section we mention the outcome of using different laser polarizations in otherwise similar simulations.

2.2.1 Preamble - Early Time Evolution of and Absorption by an Initially Cold Plasma

As was just stated, we will begin by examining the early time interaction of a plane wave with a flat surface of initially cold plasma of many times critical density. Although this geometry admittedly is a great oversimplification and idealization of what one would expect to see in a laboratory plasma (especially a compressed hydrogen plasma), we will find the results we obtain will be key to understanding the mechanisms present in a realistic interaction. Indeed, one of the most critical intuitions to take away from these numerical experiments is that, despite the fact that the initial geometry is quasi-1D, significant absorption is possible in this geometry, and occurs at quite early times in the interaction. To repeat our point, this is in direct contrast to the results of purely 1D simulations.

The physical and numerical parameters for the simulation we will look at first are as follows. The spatial geometry is again 2D, with square cells of width $0.05c/\omega_0$. Momentum and field vectors are allowed in all three dimensions, though in most situations only two electric, one magnetic, and two momentum vectors will be non-zero. The simulation box is

$25 \times 2\pi c/\omega_0$ in the x_1 direction ('long'), and $3 \times 2\pi c/\omega_0$ in the x_2 direction ('wide'). A plane-wave laser, with frequency ω_0 , is launched from the leftmost boundary of the box, starting at zero intensity at $t = 0$ and rising to a normalized intensity of $a_0 = 6$ at $t = 2 \times 2\pi/\omega_0$. The leftmost side of the box has $2\pi c/\omega_0$ of vacuum, and the plasma then extends $24 \times 2\pi c/\omega_0$ to the wall on the right. We are considering normal incidence, and the vacuum plasma interface rises rapidly. The plasma has an average charge density of $100 \times \omega_0^2 m_e / 4\pi e^2$, but with a 3% sinusoidal variation in the longitudinal direction with wavelength $4\pi c/\omega_0$ ¹³; assuming a laser wavelength of $1.053\mu m$ this corresponds to a density of $1.006 \times 10^{23}/cm^3$. The particles are initialized with precisely zero momentum in all three directions. Both electrons and ions are simulated, having an ion-to-electron mass ration of $m_i/m_e = 3643.8$. There are sixteen particles per cell for both species; the particles have cubic particle shapes (i.e. third-order splines) for both current deposit and field interpolation. Boundary conditions are periodic in the transverse direction for both particles and fields, and are absorbing for fields and reflecting for particles in the longitudinal direction.

Phase space plots of p_1 vs. x_1 at four sequential times, spaced approximately 14fs apart, are shown in figure 2.9. They are zoomed in in x_1 to a region of about a third of the entire box. The phase space density is obtained by averaging over the other phase space dimensions. Figure 2.9.a corresponds to a time only a few laser cycles after the laser has first impinged on the plasma. Although a response of the plasma is already visible, the density and energy of the accelerated electrons are quite low; the momentum of the highest energy particles is about $1 m_e c$. Though difficult to discern at this zoom level, the surface has heated up somewhat; the initially perfectly cold plasma can be seen at the rear of the target.

In the three next frames, two key changes can be seen taking place. First, the quantity of accelerated electrons and their peak momentum are rapidly increasing. In the last frame, which is taken ~ 43 fs after the laser light at the surface reached peak intensity, the highest momentum electrons are over $3 MeV/c$. Second, there is a buildup of a thin layer of hot

¹³ The reason for this modification to the density profile is that in a uniform density plasma, the electron bunches ('arms') that are sent forward can excite wakefields, whose wavelength matches to the bunch spacing if the density is a perfect square (e.g. $100n_c$). These can lead to spurious resonant interactions, which are removed by including variations like this sinusoidal one.




FIGURE 2.9: Phase space plots of $p_1 x_1$ for a 2D mobile ion simulation; (a), (b), (c), and (d) correspond to times $3.2fs$, $16.6fs$, $29.9fs$, and $43.2fs$ after laser intensity reaches its peak value at the surface, respectively. The maximum momentum of the accelerated electrons increases rapidly.

plasma at the material surface, which is discernible as a widening bulge of both forward and reverse momentum. By the last frame, the surface layer has a temperature of roughly 100 keV.

Figure 2.10 shows three more plots coincident with the last frame in figure 2.9, i.e. $43fs$ after full intensity was reached at the surface. All three frames show heat flux diagnostics. Plot 2.10.a shows the total forward heat flux being carried in to the plasma by electrons, as a function of forward position and the relativistic γ factor. It shows the same half-wavelength spaced bunches as seen in the $p_1 x_1$ plots, and shows that the energy is being carried more or

less equally by particles with γ up to 10 or more (that is to say, the color of each bunch is constant along its height). The color map scale here is linear, as opposed to the logarithmic scale used in figure 2.9. Although the high γ particles makeup a population orders-of-magnitude less dense than the total plasma density, the heat flux dynamics are dominated by their contribution. Plot 2.10.b shows the total forward heat flux versus x_1 alone (i.e. the data in fig. 2.10.a is integrated over γ), with the y-axis normalized against the average power of the incident laser. Here, we can see more clearly that the transported energy is concentrated very narrowly in sharp beams separated by twice the laser frequency (each bunch moves at c). This is in contrast to the form of the instantaneous (not time-averaged) Poynting flux of a plane wave, which would be $\cos(2k_0x_1) + 1$. The absorption can also be seen to be rapidly increasing in time (the heat flux is smaller for the bunches deeper in the plasma which were generated earlier in time) as the surface heats, increasing five-fold in the eight laser half-periods shown¹⁴.

Lastly, in figure 2.10.c the total heat flux (not just forward) versus forward position and angle of travel in radians is shown. Here we see that each of the half-laser-wavelength spaced bunches are traveling alternatively above or below the p_θ axis, by about 0.30 radians in either case; this gives a divergence at the source of roughly 35° . The reason there are no energetic electrons moving along the p_θ axis will be clear shortly.

2.2.2 Particle Tracking in Semi-1D - Acceleration Mechanism at Sharp Surfaces

Having shown from fully self-consistent 2D modeling that acceleration is possible from an initially cold, flat surface, after only a few 10's of femtoseconds, we now return to the simplified case of a preheated plasma with fixed ions and use particle tracking to determine the exact mechanism by which the energetic (high gamma) particles are being accelerated. We will see that the standing wave structure is key to this mechanism, and in particular that the structure of the magnetic field dictates how particles are, or are not, accelerated.

To study this, we will use the following simulation setup. The overall geometry is 2^{1/2}D,

¹⁴ Note that the absorption has not in fact plateaued; that the last bunch appears to be carrying less energy is because particles are still being accelerated into it.




FIGURE 2.10: *Heat flux diagnostics at 43.2 fs after peak laser intensity for a 2D simulation with mobile ions. Frame a) shows forward heat flux vs. forward position and kinetic energy; frame b) shows forward heat flux vs. forward position alone (i.e. the data from (a), integrated along the vertical axis). Frame c) shows total heat flux vs. forward position and angle of propagation. The total absorbed energy can be seen increasing rapidly in b) (note the front peak is still in the process of gaining energy), and a) shows that the heat flux is carried evenly by particles of all energies (indicated by each bunch being roughly the same color throughout.) The accelerated bunches alternate in having transverse momentum one way then the other, as seen in c).*

with a transverse width for the simulation box of $10c/\omega_0$, and a longitudinal length of $25.2c/\omega_0$. The front of the box is a vacuum region with a length of $3.15c/\omega_0$, the rest of the box is filled with a $100n_c$ density electron plasma¹⁵, and the ions have infinite mass. The electrons are initialized with an isotropic thermal distribution with a temperature of 153 keV ($\gamma_{th}\beta_{th} = 0.548$). A plane-wave laser with a normalized vector potential $a_0 = 6$ is launched from the left-hand wall towards the plasma, with $\vec{k}$ parallel to the surface normal. The laser rise-time is less than a laser period, after which it remains constant. The cell size is $\frac{1}{100}c/\omega_0$ in both directions (10 cells per skin depth), there are 16 particles per cell, and the particle shapes are cubic splines. The transverse boundary conditions are periodic, and the longitudinal boundary conditions are absorbing for fields and thermalizing for particles.

As stated, we will be using 2D OSIRIS for these simulations; however we should note that – due to the fixed ions, the short time-frame of these simulations, and the translational invariance in the second direction in the initial setup (plane wave, uniform density, uniform temperature, and periodic boundary conditions) – in this case much of the physics will be 1D and we could perhaps term this a semi-1D acceleration mechanism. However we still wish to emphasize that the overall acceleration mechanism is inherently two-dimensional, as it is only made possible by multi-dimensional heating.

Figure 2.11 shows phase space (x_1 vs. p_1 and p_2 vs. p_1) results for the simulation. In fig. 2.11.a, which shows the particle density versus forward momentum and forward position, as expected we see a result similar to the 1D hot result from an earlier section, with hook-like arms of particles streaming into the plasma with a spacing of half a laser wave-length. Fig. 2.11.b shows the forward and transverse momentum, and here we can see again of course the forward jets of electrons; however it's quite noteworthy that there are no particles moving straight forward ($p_2 = 0$), rather there are two sets of jets with finite either negative or positive p_2 .

To elucidate the acceleration mechanism we will use particle tracking. Particle tracking, as the term implies, involves tracking the trajectories of individual particles throughout the duration of the simulation. This permits following their exact movement through $\vec{p}$ - and $\vec{x}$ -

¹⁵ A gentle density ramp is again added to the plasma, to avoid resonant effects.




FIGURE 2.11: *Phase spaces for a hot, semi-1D simulation. a) Forward momentum vs. forward position. b) Forward and transverse momentum space. Particle positions for tracked particles are indicated by black dots. The transverse motion of the alternating bunches is reflected in the two peaks in the $p_1 p_2$ phase space, on either side of the p_1 axis.*

space and to record the fields which they “see” during that movement. In general we do not know which particles will be ‘interesting’ before the fact; therefore it is usually necessary to run a full simulation once and then select the particles which display the interesting features at the end of the simulation. We then have to run the exact same simulation again, this time recording their tracks. It is of course therefore necessary to re-run the simulation with bit-wise equivalent results (including all initial conditions, so in particular the same pseudo-random number generator seeds need to be used everywhere), and the ability to do this even when working with billions or even trillions of particles spread across thousands of processors is a very useful feature of the OSIRIS code.

The particles which we will select to track are shown as black dots in figure 2.11. Two groups of electrons have been selected. One group is all within one of the 2ω jets, with a range of energies spanning most of the height of the jet. The second group consists of a smaller number of particles which are also traveling with the jet, but which have energies consistent with the thermal plasma; these were selected with the hope of seeing how they failed to gain energy.

Tracks for all of these particles are shown in the bottom row of figure 2.12. Tracks are shown in position space. The tracks start at the initial positions and end at the marked times. The color of the tracks indicates the particle energies. The current position is shown as a black dot. Three tracks of interest to be discussed later are identified with circles around their end points. All particles start and remain near the vacuum-plasma boundary and so only this region is shown. Frames are shown for three different times: just as the accelerated particles leave the plasma region ($t\omega_0 = 12.25$), roughly the time of maximum excursion in to the vacuum region for all such particles ($t\omega_0 = 13.25$), and the point in time when most particles are re-entering the plasma ($t\omega_0 = 14.25$). The top row of the figure shows line outs of the transverse electric and magnetic fields (‘laser fields’) at the same times as the tracks, along with the longitudinal electric field (‘electrostatic field’).

2.2.2.1 Movement of the particles

Several observations are immediately apparent. In the first frame, most electrons are beginning to cross the plasma-vacuum interface and leave the plasma. All of these electron’s initial motion was upward, i.e. they all started with a positive p_2 (and a negative p_1). As these electrons reached the interface they then rotate such that they are moving straight to the left. A typical electron from this group is indicated by the blue circle. Showing a different behavior, a few electrons – one such indicated by a red circle – have already turned around, and each of these approached the interface at near a normal incidence line ($p_2 \sim 0$). In the second frame one sees that after having traveled a short distance straight into the vacuum, particles are then accelerated upward (i.e. they are gaining energy). Finally, in the third frame, the particles can be seen to be rotated so that they move back into the plasma, gaining at most a small, final bit of energy as they cross the vacuum, before re-entering with a component of velocity still traveling upward.

The final energy of the particles is correlated with the distance they penetrate into the vacuum. This is most clearly seen in the track highlighted with the black circle. As can be seen, this electron travels farther into the vacuum region than any other particle, and is also




FIGURE 2.12: Bottom row: Tracks of accelerated and non-accelerated particles shown in their movement in position space, at three time frames. Particle energy is indicated by track color. Key tracks are identified by circles, as explained in the text. Top row: Line-outs of the laser E and B field and longitudinal E field, at the same time as the particle track frames. Particles can be seen to be rotated by the laser B field before and after being accelerated by the laser E field.

the only track shown which attains an energy of $\sim 5\text{MeV}$ (indicated by the red color of the track). Although not shown here, this particle actually *loses* some energy at very end of its travel through the vacuum. Despite this it is still the most energetic particle at the end of the acceleration phase. Also of note is that, owing to the greater distance this electron travels out into the vacuum, it ends up re-entering the plasma at a later time than all other electrons; this point will be returned to later.

Contrasting this behavior, we have the movement of the particles which are not accelerated, one of which as noted earlier is highlighted with a red circle. These particles are seen to travel towards the surface almost directly or with a slight amount of downward momentum. They execute the same Larmor rotation as the other particles when they reach the surface; however in their case this motion rotates them back into the plasma rather than out of it. They end the interaction traveling back towards where they came without having gained any significant energy at any time.

2.2.2.2 Structure of the fields

In the row above the particle tracks we show line-outs of the three electromagnetic fields¹⁶, the laser fields E_2 and B_3 , and the electrostatic field E_1 . There are a number of important facts about the interaction which we can discern from this time series of the fields, some of which are intuitive or at least straight forward and some of which are slightly subtle. The first is that the laser fields are very close to those of a perfect standing wave as one would expect from a finite, collisionless reflector. The form of such a field is trivial to derive (if the effect of the magnetic field on the particle motion is ignored as can be done in the linear regime) so we will not repeat the work here; but we will remind the reader that in a standing wave the electric and magnetic fields are one quarter period out of phase in both time and space, with the boundary conditions on the fields and their derivatives giving a maximum

¹⁶ Although the simulation monitors all three components of both the electric and magnetic fields, as one would expect, it can be easily verified from the time-dependent Maxwell's equations that in two-dimensions the two in-plane components of the electric field and the out-of-plane magnetic field, along with the in-plane momentum components, form a closed system. The other three components will be zero, or at most thermal fluctuations, and so can be ignored.

or anti-node of the B field at the material surface, and hence zero or node for the electric field at the same location.

However the overdense plasma does not act like a perfect, collisionless reflector, in three noteworthy ways. First, intuitively we know the reflection factor isn't unity because particles are being accelerated. However this effect probably has a small impact on either the fields or tracks of individual particles, as can be inferred from the fact that the peak field values are quite close to the perfectly reflecting value of $2a_0$, consistent with the high reflection coefficient seen in all but the highest temperature 1D cases. And furthermore the results of numerical calculations of the theoretical fields in high absorption regimes (modeling the absorption, unphysically, as collisional - the motion of the electrons within the fields is quite collisionless) show both the field structure and the particle trajectories to surprisingly insensitive to this parameter; this will be shown in more detail in the next section. Still, a careful inspection of the figures does show some distortion which is likely due to the finite absorption; most noticeably the zero-crossing of the fields don't remain perfectly stationary in space. That this could have a small effect on the particle motion shouldn't be discounted.

Second, we are in the regime of $a_0 > 1$, meaning we in fact cannot fully neglect the magnetic field in the response of the material. Rather, the laser B field results in a net force inward on the plasma surface, a well known effect which is often erroneously termed a ponderomotive force but actually is a direct (not time averaged) force from the laser during each period. The result of this force is that the electrons are pushed back across the ions until an electrostatic surface field counterbalancing the $\vec{v} \times \vec{B}$ pressure exists (during each laser cycle). The effect of this field on a test particle therefore would be to eject it out into the vacuum, and hence might be expected to be related to injecting electrons into the laser field. Such a surface field can indeed be seen in the line-outs. However only in the last frame is the sign of the field consistent with electron ejection (a force to the left for the electrons); the previous two times show fields of comparable strength but in a direction which would push the electrons back into the plasma, not pull them out into the vacuum.

The reason for the oscillatory nature of the sign of the electrostatic field is related to the third characteristic in which the plasma differs from the linear analysis, which is that the

plasma surface (or in this simplified case the whole plasma) is hot enough that the thermal effects cannot be neglected. The effect of the finite temperature of the electrons is that it will cause them, in the absence of the laser field, to attempt to diffuse out into the vacuum, causing a charge separation which creates an ambipolar field pulling the electrons back to the plasma. In the presence of the laser, the ambipolar field opposes the light pressure force, and the electrostatic field is able to switch between either sign.

It is also insightful to compare the phases of the fields with that of the accelerating motion of the electrons. First, upon examining the magnetic field, it can be seen that particles which gain significant energy approach the surface just as the field is decreasing in time towards zero, and then travel into the vacuum region as it changes sign. As they penetrate into the vacuum they gain energy and are then rotated back towards the plasma. They enter the plasma and hence exit the field just as it is reaching a new extremum (in this case minimum). It is clear from the tracks that in doing so, the electrons use the strong but decreasing magnetic field to rotate out towards the vacuum (they leave the plasma nearly normal to the surface), then later use the opposite phase of the field to rotate back into the plasma.

Next, looking at the laser electric field E_2 we see the electrons enter the field just as it starts increasing from zero, and exit the fields (enter the plasma) just as it passes through zero and reverses sign. The electrons are thus able to utilize nearly the entire laser half-cycle to accelerate, which naturally maximizes their energy gain. Obviously, the electrons gains the most energy if it can sample the maximum of the electric field. In the standing wave pattern $|E_2|$ is a maximum a half wavelength into the vacuum, and this explains the need of the accelerated electrons to be rotated out into the vacuum by the magnetic field.

Lastly, we can look at the behavior of the electrostatic field, where we can see that it is out of phase with the motion of the accelerated electrons, pushing them back into the plasma as they exit and pushing them out towards the vacuum as they re-enter the plasma. However this field is much weaker and it exists in a narrower region of x_1 than the other two fields; therefore its effect can be considered minimal.

2.2.2.3 Acceleration mechanism

We are now in a position to connect the motion of the accelerated particles with the field structure, and to draw conclusions about the phenomenology of the acceleration mechanism. First, we notice that there is no longitudinal electric field of significant strength, and therefore that a particle's best bet for strong acceleration is to somehow interact directly with the transverse E field (magnetic fields, of course, doing no work.) We also can conclude from the weak E_1 that there is no force which will push the particles out into the laser field, and therefore they must do so 'themselves,' i.e. their thermal motion must let them leave. This explains the absence of absorption in cold plasmas. Furthermore, we notice that the distance from the plasma surface to the region of greatest energy gain is one quarter of the laser wavelength, and if we assume that the particles must complete their interaction with the light field before it switches signs (that is within a half laser period, otherwise their motion would almost certainly become chaotic, among other issues) we therefore find that the particles must be moving near the speed of light (they need to move out and back a quarter wavelength every half laser period). This explains why the plasma must not just be warm but must in fact be relativistically hot for electrons to gain the maximum energy.

Next, we come back to the fact that in a standing wave the transverse electric field has a node and the magnetic field has an anti-node at the surface. Therefore, in order for an electron to reach the anti-node (maximum) of the electric field it must cross through the magnetic field. Furthermore the phases between the fields are such that if the electron exist the plasma when the magnetic field is zero then the electric field is beginning to decrease. As a result, by the time the electron gets a quarter wavelength into the vacuum the electric field has vanished. Therefore, the best strategy for an electron which seeks to gain energy would be to use the surface magnetic field to its advantage by approaching the surface obliquely and then being rotated out into the vacuum normal to the surface by the magnetic field. If it does this while the magnetic field is decreasing then the electric field would be increasing such that it is large when the electron gets a quarter wavelength into the vacuum. This is indeed what is seen in the tracks and it explains the observation seen in one-dimension, that

the electrons need large transverse momentum in order for laser absorption to occur.

We also note that because the electrons which gain the most energy must also be the one which penetrated deepest into the vacuum, these electrons must also be the last to re-enter the plasma. This was already seen from the particle tracks of the most energetic electrons. The consequence of this is that lower energy (but still relativistic) particles will be traveling in front of the higher gamma ones as they move through the plasma. This is in contrast with the behavior of non-relativistic beams generated in cold plasmas, where the high energy particles out pace all others and end up leading the beam. For the 2D (or hot) plasma cases the phase space shape of the beam comes not from the kinematics of its propagation, but rather is a frozen in feature of the acceleration mechanism itself.

Now have explained why all the accelerated particles must be initially traveling transversely (in addition to longitudinally that is, since otherwise they wouldn't be approaching the boundary), we can go back and explain an interesting feature we noticed in 2.10.c and 2.11.b, which is that there are no accelerated particles traveling straight into the plasma ($p_2 \sim 0$). The reason for this is the conservation of transverse canonical momentum in the semi-1D geometry (cf. appendix A). Electrons start inside the plasma where the vector potential from the laser is zero (the plasma is over-dense). Therefore, when they return to the plasma they must have the same transverse momentum they are initialized with, no matter how complicated their trajectories outside the plasma. Therefore, since only particles with a non-zero p_2 get accelerated, then upon returning to the plasma they all must have non-zero p_2 . Therefore, the bifurcation seen in the p_1 vs. p_2 phase space is indeed a proof that the acceleration mechanism “selects” for electrons with an initially high transverse momentum.

2.2.2.4 Tracks in phase space

By examining in detail the particle tracks in phase space, we can further see the characteristics of the accelerating motion. The top row of figure 2.13 shows $p_1 x_1$ and $p_2 p_1$ phase space, along with the tracks of six particles with a range of final energies (note that these are not the same particles tracked earlier in figures 2.11 and 2.12). The numbers in the figures show




FIGURE 2.13: Phase space plots for an OSIRIS simulation of laser acceleration at a sharp surface with a hot plasma (top row), and for a simulation for the same temperature plasma but using a simpler code ('single particle' code) where the fields are those of a perfectly reflected laser, i.e. the particles have no effect on the fields (bottom row.) The OSIRIS figures also have the tracks for a few selected particles superimposed, along with numbers indicating the correspondence of the electron bunches between the two phase spaces. Although the physics is greatly simplified in the single particle code, the phase space plots look very similar to the PIC results.

the correspondence of the jets between the two phase spaces, showing how each alternating jet in (a) is traveling into or out of the page (i.e. in the x_2 direction), in addition to into the plasma (in the x_1 direction).

In figure 2.13a we see all the particle tracks initially traveling towards the surface, rotate out into the vacuum, (a few of the tracks seem to spend a bit of time twisting around in the skin-depth B field before managing to enter the standing wave in correct phase) sharply gain a large amount of forward momentum in a thin layer at the surface, then finally penetrate back into the plasma and travel roughly ballistically. The nesting of the tracks clearly shows the correlation between distance into vacuum and total energy gain, and the origin of the backward-curving hook shape.

The interplay between transverse and forward momentum during the acceleration phase can be appreciated by looking at figure 2.13b. Particles can be seen being greatly accelerated transversely, whipping out into p_2 , and then having all that momentum gain rotated back into the direction normal to the plasma surface. A jet can also be seen in mid-acceleration on the left-hand side of this figure, and correspondingly at the surface in the p_1x_1 phase space.

2.2.2.5 Effect of the electrostatic field

We return for a moment to emphasize an assertion we made earlier, that the longitudinal, ‘electrostatic,’ field does not play a crucial or in fact in any way significant role in the acceleration of the particles. The reason we wish to revisit this point is that many of the acceleration or heating mechanisms proposed by others [32, 45, 44] in similar geometries but different parameters regimes, along with other absorption mechanisms proposed for the same parameter space [81, 47], rely on the interplay of an electrostatic field either directly or indirectly to accelerate the particles. This is a critical point because by relying on a longitudinal field to mediate the energy exchange, the necessary role of high temperature and transverse momentum is missed.

We postulated earlier that the electrostatic field had little or no effect because we had observed that it was weak, thin, and in fact of oscillating sign. To demonstrate that this is

true we use a diagnostic that separates the total work, $\int \mathbf{v} \cdot \mathbf{E} dt$, into the work done, on the same tracked particles as those in figure 2.13, by the transverse and longitudinal fields, i.e. $v_2 E_2$ and $v_1 E_1$, respectively. These results can be seen in figure 2.14, which is centered around the window in time where there is energy gain. In figure 2.14.b energy can be seen being smoothly gained from the laser field, (except for the particle which was out of phase with the surface magnetic field and hence even lost energy). In contrast, the electrostatic field has almost no effect on the energy – and what effect it does have is as likely to be decelerating as accelerating – as can be seen in figure 2.14.a. The total energy is shown in 2.14.c, which as expected shows the total acceleration and the transverse acceleration being almost equal; a small electrostatic deceleration as the particles re-enter the plasma can be more clearly seen here as well.

One of the reasons the E_1 is not a major factor in the acceleration or heating is that the two effects which create the field oppose each other, i.e. the radiation pressure pushes the electrons inward and the thermal expansion pulls them outward. The resulting surface field is an ambipolar field which attempts to counteract both the radiation pressure and the thermal pressure. However, although we do not show particle tracks for the cold case, we note that even in this case where no such cancelation occurs and the radiation pressure leads to a surface field which would attempt to eject particles into the vacuum, no acceleration is seen. Simply put, we find that in the high density, high laser-intensity regime, the electrostatic field is too weak compared to the laser fields to either be a direct (actually doing work on the particles) or indirect (i.e. injecting electrons into the vacuum) source of acceleration.

Lastly we want to emphasize that although the electrostatic field at the surface does not have an important effect on the laser absorption or particle motion, that is not to say that longitudinal fields do not have an effect on the particle energy dynamics. Specifically, the wake fields of the electron bunches inside of the plasma have the effect of mixing the energy balance between the high- and lower-gamma particles. This can be seen in figure 2.14.d, which shows the integral of $\mathbf{v}_1 \cdot \mathbf{E}_1$ over the first 12 fs or 4 μm they spend inside the plasma, where exchange of energies of a few times the rest mass are already apparent. This effect is discussed in more detail in chapter 4, though there in the context of macro-particle stopping,




FIGURE 2.14: The integral of the dot product of particle velocity and the electric field, as a function of time and broken down by components, shows minimal energy gained due to the longitudinal field (a), and significant energy gain due to the transverse field (b). In (c) we show total energy as a function of time, which mostly tracks (b) but also more clearly shows the small deceleration for some particles as they enter the plasma. In (d) the effect of the longitudinal field as the particles travel through the plasma is plotted, which demonstrates that electrons both gain and lose energy as they interact with wakefields. All tracks are colored by particle energy

an artificial effect.

2.2.3 Single Particle Model and Numerical Calculations

We have shown that the field structure in the fully self-consistent PIC model is quite close to that from perfect reflection at the plasma surface; and we have seen from the fields and from particle tracks that it appears that it is the thermal motion of the particles that causes them to leave the plasma and interact with the laser fields in vacuum. To further study this point, we turn now to numerical calculations of particle trajectories through imposed fields, to attempt to see what is the least set of assumptions we need to reproduce the PIC results.

Specifically, we will use what we call a ‘single particle model’ or ‘particle tracking’ code, meaning that we use electric and magnetic fields calculated analytically before hand, and we advance individual test particle trajectories through these fields. Although we call this model ‘single particle,’ and that accurately describes the underlying calculation, in practice we will actually be initializing an ensemble of particles and advancing them simultaneously through the fields. This facilitates comparison with the PIC phase space results. The code itself is written in Python and Fortran, and uses the same second order accurate, relativistically correct, time centered Boris particle pusher as is in OSIRIS.

More specifically, the ensemble we will be using will fill a volume in $x_1 p_1 p_2$ phase space on a regular grid. The particles however will have a ‘weight’ the based on the value of a Maxwellian at that point in momentum space (the weights are used only for the diagnostics); the temperature used for weighting will be $120keV$ unless otherwise specified. The laser will be incident from the negative x_1 direction and will reflect at the origin, and so we will initialize the plasma from $x_1 = 0$ up to some finite x_1 . The value of this only matters in as far as there are still particles traveling towards the interaction region at the time at which we analyze the diagnostics, which is done after a quasi-steady state is reached; four times the laser wavelength is typical. Since only particles traveling towards the surface will ever interact, we only initialize the half of phase space with negative p_1 ; otherwise all velocities up to roughly seven times the thermal momentum are considered.

As stated, the fields used in the single particle code are those of reflection at normal incidence from a cold finite plasma with a step function density gradient at $x_1 = 0$. The plasma response is assumed to be linear, i.e. we ignore the $\mathbf{v} \times \mathbf{B}$ force. This is reasonable for intense lasers if they cannot penetrate into the plasma. Time, length, wavenumber, plasma frequency, and vector potential are all normalized to be unit-less by the substitutions $t\omega_0 \mapsto t$, $xk_0 \mapsto x$, $k/k_0 \mapsto k$, $\omega_p/\omega_0 \mapsto \omega_p$, and $eA/m_e c \mapsto A$, where e is charge of electron, m_e its mass, c the speed of light in vacuum, ω_0 is the frequency of the laser and $k_0 = c/\omega_0$. The normalized complex vector potential for such fields is then given by

$$\mathbf{A}(\mathbf{x}, t) = A_2(x_1, t)\hat{\mathbf{x}}_2 = \begin{cases} a_0(e^{i(x_1-t)} + Re^{i-(x_1-t)}) & x_1 < 0 \\ a_0Te^{i(kx_1-t)} & x_1 > 0 \end{cases} \quad (2.3)$$

where a_0 is the magnitude of the vector potential of the incident wave, k is the wavenumber inside the plasma, and the transmission T and reflection R coefficients are given by

$$R = \frac{1 - k}{1 + k} \quad (2.4)$$

$$T = \frac{2}{1 + k} \quad (2.5)$$

It should be understood that it is the real component of the complex vector potential and fields which represents the physical quantities.

The normalized dispersion relation for an electromagnetic wave in a plasma is given by $k^2 = 1 - \omega_p^2$. From this it can be seen that when the wavenumber is k purely real (and strictly less than 1 in this normalization) it corresponds to a collisionless plasma with plasma frequency ω_p less than the frequency of the incident wave; and that when k is purely imaginary, it corresponds to a collisionless plasma with plasma frequency greater than that of the incident wave (i.e. $k = \sqrt{1 - \omega_p^2} = i\sqrt{\omega_p^2 - 1}$). A real wavenumber, k , therefore corresponds to transmission into an underdense plasma, and is not pertinent to our question, and so we will not consider this case. However, it can be quickly verified from eq. 2.4 that the magnitude of the reflected wave is unity for all k with zero real component, and so to study the effect of finite absorption within this model we must consider a complex $k = k_r + ik_i$. Doing so is mathematically equivalent to introducing an Ohmic drag term $\eta\mathbf{J}$ to the equation

of motion of the fluid, $\mathbf{E} = 4\pi/\omega_p^2 d\mathbf{J}/dt + \eta\mathbf{J}$. However it bears emphasizing that any realistic collisionality in the regime of interest would lead to minimal resistivity and hence minimal absorption, and so this should be considered a method to most expediently add absorption to the linear model and not strictly speaking to model the correct physics¹⁷.

The first and arguably most important result we will look at is also the simplest (important, that is, because it's so simple to analyze), where we allow $k \rightarrow i\infty$, which is of course the case of a perfect reflector. The vector potential is then given by

$$\mathbf{A}(\mathbf{x}, t) = A_2(x_1, t)\hat{\mathbf{x}}_2 = \begin{cases} a_0(e^{i(x_1-t)} - e^{i(-x_1-t)}) & x_1 < 0 \\ 0 & x_1 > 0 \end{cases}$$

and the fields, found from $\mathbf{E} = -\partial\mathbf{A}/\partial t$ and $\mathbf{B} = \nabla \times \mathbf{A}$, are simply

$$\begin{aligned} \mathbf{E}(\mathbf{x}) = E_2(x_1, t)\hat{\mathbf{x}}_2 &= \begin{cases} a_0(\sin(-x_1 + t) - \sin(x_1 + t)) = 2a_0\sin(x_1)\sin(t) & x_1 < 0 \\ 0 & x_1 > 0 \end{cases} \\ \mathbf{B}(\mathbf{x}) = B_3(x_1, t)\hat{\mathbf{x}}_3 &= \begin{cases} a_0(\sin(-x_1 + t) + \sin(x_1 + t)) = 2a_0\cos(x_1)\cos(t) & x_1 < 0 \\ 0 & x_1 > 0 \end{cases} \end{aligned}$$

The results are shown in the bottom row of figure 2.13, aligned with the self-consistent PIC results at the same phase of the interaction for ease of comparison. As can be plainly seen, this rudimentary model has already captured almost all the salient features of the electron acceleration. This clearly demonstrates the importance of the standing wave structure for the acceleration of particles, and the lack of a need to include any electrostatic field to explain either the acceleration or the injection into the field.

Next we investigate what effect the inclusion of a skin depth layer will have on the acceleration of particles. The first row of figure 2.15 shows phase space plots for $k = 10i$, which corresponds to the same density as we have been using in the OSIRIS simulations (in fact $n = 101n_c$, but the difference should be trivial). The next three rows show results for $k = 3i$, $2i$, and i , corresponding to $n = 10n_c$, $5n_c$, and $2n_c$ respectively. As can be seen,

¹⁷ The algebra of the relationship between η and k_r is trivial but complicated, and so since this is a non-physical scenario we will not provide the detailed algebra here.




FIGURE 2.15: Phase space plots for the single particle code with finite skin depth; the value of the wavenumber in the plasma is shown to the right of the figures. Only when the skin depth becomes on the order of the laser wavelength do the results start to significantly differ from perfect reflection.

the inclusion of a realistic skin depth leads to similar phase space plots and hence does not change the acceleration process significantly. This result is consistent with the PIC result (which would also have a skin depth effect), and so could have been expected. It is not surprising that the inclusion of the skin depth does not alter most of the acceleration process since it is very short compared to the laser wavelength.

What is rather surprising is that even the inclusion of unrealistic skin depths doesn't modify the phase spaces very much. For a density of ten times critical, the differences in the particle behavior are almost imperceivable, and even at $5n_c$ the apparent differences are only quite subtle ones. It is not until we get down to the definitely unrealistic density of twice critical - unrealistic, among other reasons, because the plasma will be relativistically transparent to any pulse with a_0 much greater than unity, and certainly to the $a_0 = 6$ shown here - that we start seeing the plasma response be altered to a noteworthy extent.

We can also use this model to investigate, roughly, the effect of laser absorption on the phase space of the electrons. The absorption process requires a standing wave. If the absorption is too high then the reflected wave is less than the incident wave and the standing wave structure is modified. It is therefore interesting to examine how the acceleration process is modified as absorption increases. As mentioned earlier, we introduce absorption by using a complex k , which is equivalent to introducing an effective resistivity to the plasma. Again this is not the correct physics for laser absorption in this regime; rather the true absorption is due to the disruptive movement of the electrons forward into the plasma at every half laser cycle (due to the neglected magnetic force), thereby preventing them from harmonically cycling their energy back into the field. And so, these results should not be taken too literally or be used to compare too directly with quantitatively equivalent absorption rates. Still we think the exercise is somewhat instructive, and so we include it here.

In figure 2.16 we show the same phase spaces as figure 2.15, for four effective densities and for the maximum absorption at that density. The highest possible absorption as a function of k_i is monotonically increasing in decreasing k_i (the closer k_i is to unity, the closer the plasma is to vacuum, the deeper the wave can travel into the plasma, and so the longer distance




FIGURE 2.16: Phase space plots for the single particle code with finite skin depth and non-zero laser absorption, modeled by having complex wavenumber (shown on the right) in the plasma (i.e. the plasma is modeled as having some collisionality). The first and second rows from the top (where absorption is 9.5% and 18%, respectively) have similar phase space dynamics to absorptionless cases. The third and fourth rows (where absorption is 28% and 38%, respectively) begin to have different dynamics.

the wave has to damp over), and occurs for a given k_i at $k_r = \sqrt{k_i^2 + 1}$ ¹⁸. For the k_i plotted here, $10i$, $5i$, $3i$, and $2i$, the total absorption is approximately 9.5%, 18%, 28%, and 38%, respectively. At the lowest absorption, 9.5%, the phase space appears barely affected; even almost doubling the absorption, to 18%, doesn't change the peak acceleration energy, though it does somewhat change the shape of the $p_1 x_1$ phase space right at the peak of the jets. It's not until the absorption gets above 25%, for $k = 3i + \sqrt{10}$, that the peak energy begins to decrease, though qualitatively the 2ω jets are still clearly present. For 38% absorption the phase spaces are starting to be grossly modified, and the peak energy of the jets is further lowered.

One point we can unambiguously conclude from these calculations therefore is that, even in the presence of absorption of up to a quarter of the incoming light, the basic acceleration mechanism remains unchanged from that of perfect reflection at the surface. When absorptions in our model go above this, this no longer appears to be the case, even when comparing results of calculations with similar skin depths (c.f. especially figs. 2.15.f and 2.16.h). However nonetheless such skin depths are only consistent with much lower densities than we used in our PIC simulations or would be expected in an experiment, and in using them we modify the field structure in ways significantly beyond just the effects of the absorption. In addition, absorption ratios this high – at least in one-dimension – were also associated with extremely high temperatures, even higher than we have here, and so the reflection physics would likely be modified in ways not captured here (from the thermal expansion of the plasma surface, for instance.) Therefore we conclude that for higher absorptions we've really stretched this model beyond its validity, and cannot from it determine conclusively how absorption would be affected in a realistic scenario. We will have to rely on PIC simulations for data in that regime.

At the risk of belaboring the point, we include one more calculation with our single particle code demonstrating the robustness of the acceleration mechanism to realistic skin-depth

¹⁸ This relation is readily found by taking the magnitude squared of the reflectivity coefficient from eq. 2.4, setting the derivative with respect to k_r to zero, and solving for k_r as a function of k_i . Assuming k_r is given by the above formula, the total absorption A can be calculated again from eq. 2.4, is given by the (not terribly illuminating) equation $A(k_i) = 1 - \sqrt{(2 + k_i^2 - 2\sqrt{1 + k_i^2})/k_i^2}$

effects. Specifically we wish to show that particles which lack sufficient initial momentum are excluded from the acceleration mechanism, a point we already saw in one-dimensional PIC simulations and which was implicit in the phase space densities in figure 2.15 but which we now calculate more explicitly. We think this is important because it demonstrates the significance of symmetry and canonical momentum conservation, which is a point of some confusion in the literature. More specifically, we wish to compare our results with those of G.E. Kemp et al. 2013 [76], who claim to have found that by including a realistic skin depth initially cold particles are able to accelerate in the wave.

Kemp et al. use a similar calculation to our single particle code to arrive at this conclusion, with one important distinction in the way in which they set-up the problem: instead of initializing particles with a range of initial positions all at some initial time, they initialize all of their particles at the plasma surface but at different phases of the incoming laser. This method has the advantage of more explicitly sampling the entire range of phases of the wave; but it does have some disadvantages too, for instance the fluence of particles traveling normally and obliquely to the surface are treated as equal. It can however have some validity, if treated carefully – but it can also lead to some issues if not treated carefully, as we will see below.

The results published in [76] are reproduced in the top row of figure 2.17, and results using our code modified to use their initialization and diagnostic are shown in the bottom row of the same. The parameters for the calculation are that the laser has a normalized vector potential of $a_0 = 1.5$, and that the plasma has a density of $n = 45n_c$. The diagnostic looks at the energy and angle of incidence relative to the surface normal for accelerated particles when they re-enter the plasma, with angle as the abscissa and energy as the ordinate. The color map shows the initial energy of the particle (direction of the initial velocity vector is not tracked), with lower-energy particles layered on top of higher ones, which is to say the color represents the lowest possible initial energy necessary to exit the wave with a given final energy and direction.

Figures 2.17.ia & 2.17.a show the results for the case of an infinite plasma density, that is no skin-depth, for Kemp’s code and ours, respectively. Consistent with our various earlier




FIGURE 2.17: The effects of using the single particle code but with particles initialized at a range of times all at the surface rather than a range of positions at one time. All figures show final energy and angle of travel versus initial energy of the particles. The top row is reproduced from GE Kemp et al. [76]; (ia) is for the case of perfect reflection at the surface, and (ib) is for the case of a finite skin depth. The bottom row shows the results of similar calculations using our code; (a) & (b) are equivalent to (ia) & (ib), respectively, and (c) shows results for inclusion of a finite skin depth with a correction to the canonical momentum. The model erroneously shows acceleration of initially cold particles when the momentum is naïvely not adjusted to take into account the vector potential of the laser.

calculations, Kemp sees that only particles with high initial kinetic energies can be accelerated by the wave, and consistent with Kemp we reproduce that result using our code and their diagnostic. Figures 2.17.ib & 2.17.b shows the data for a similar calculation, where the particles are initialized at the surface with the same energies but now the fields are those of reflection from a $45n_c$ plasma, again from [76] and using our code respectively. Here now Kemp finds that particles which are initially relatively cold are in fact able to gain energy, and again we are able to reproduce that result with our single particle model.

There however is an error in this calculation which is somewhat subtle. It arises from the fact that, by initializing a particle with zero physical momentum at $t = 0$ arbitrarily inside the wave, one has not initialized it with zero canonical momentum. In essence for the particle to be at rest at that location at that time, it must have started out somewhere else with an energy which is non-zero¹⁹. To correct for this somewhat we can initialize our particles such that their canonical momentum is consistent with the physical momentum for a given energy; that is we find its ‘initial’ energy (now thought of as being at some earlier time) based on some longitudinal and transverse momentum, but then from the latter we subtract the value of the laser’s vector potential at whatever phase of the wave the particle is injected into. We are still making some assumptions that may be unphysical; most obviously the longitudinal momentum may also have been affected when the particle ‘traveled’ from where we imagine it started in to the wave. This is one reason we choose to initialize our particles with a range of x_1 ’s and have then drift into the wave instead.

It turns out in the end however that this correction is enough to approximately recover the original result, as is shown in figure 2.17.c. It in fact appears that the final energy is an even stronger function of the initial energy than was the case without any skin depth. We would argue that this result really should not be surprising, since it is in agreement with the self-consistent PIC results shown earlier. However arriving at it was probably worth while, we believe, since it nicely illustrates the importance of the underlying symmetries of the problem.

¹⁹ An alternate view would be that the particle *was* created at $t = 0$, i.e. it was injected from an ionization event. If we consider our plasma to be fully ionized at this point (which is consistent with a constant plasma density) we can exclude this possibility.

2.2.4 Initial Heating Mechanism - ‘Slipping Stream’ Instability

So far we have discussed the acceleration of electrons in an essentially one-dimensional sense. However, we have also made the point that the acceleration is dependent on the transverse temperature, and that any heating in the transverse direction is strictly forbidden in 1D due to the conservation of transverse canonical momentum; hence we have called this a 1D acceleration mechanism which will only be seen in higher dimensional simulations. We have demonstrated that in two-dimensional simulations the plasma does indeed heat sufficiently (to $\sim 100keV$) to access the acceleration regime, but we haven’t explored how exactly the plasma surface is able to gain such high temperatures.

Although at first it may be somewhat surprising, it actually isn’t very difficult for the surface to reach these temperatures. The reason has to do with the fact that, before any other mechanism for absorption or electron acceleration is possible, any energy deposition which does occur will be confined to where the laser field is non-zero, which is a layer only a skin depth thick. Although this region is quite thin heating it is still sufficient to create electrons which can then interact further with the laser; and since this region is so thin even if only a small fraction of the incident laser energy is absorbed it can still heat up to the required temperatures quite quickly. For instance, if we use as typical numbers a laser intensity of $5 \times 10^{19}W/cm^2$, a plasma density of $10^{23}/cm^3$, and a skin depth of $0.1\mu m$, we arrive at an available power of $50W/electron$; we then only need to assume a coupling efficiency of 3% for this layer to heat up to $100keV$ ($= 1.6 \times 10^{-14}J$) in 10fs. Hence, when the only absorption which is possible is confined to the surface, it would in fact be surprising if the plasma there did not heat up substantially.

It is nonetheless important to identify processes that can allow an intense laser to be absorbed by a collisionless and “cold” plasma at normal incidence to a sharp interface because many absorption mechanisms do not apply. For instance Brunel absorption [44] is not possible because of the normal incidence, resonant absorption [32] is not possible because of the sharp interface (and the normal incidence), and inverse Bremsstrahlung [25] is not possible because of the lack of collisions. Even so, based on both theory and on the results we have

seen in simulations, we believe even for normal incidence at sharp interfaces the physics is complicated and involves multiple mechanisms, often happening simultaneously. However there is one mechanism which we believe is particularly important, especially early in time, which can be described as a Kevin-Helmholtz-like instability²⁰ the effect of which is to break the transverse invariance and so allow non-conservation of transverse canonical momentum. Later in time ion motion and surface rippling occurs leading to additional heating and absorption processes.

More specifically the mechanism which we believe leads to the initial heating – which we sometimes call ‘bootstrap heating’ because it unlocks other mechanisms which can themselves lead to more heating – is related to the transverse flow of electrons inside the skin depth layer, in particular to the exponential decay and therefore inhomogeneity in the longitudinal direction. Strictly speaking such a flow is not stable (even for fixed ions), which is somewhat intuitive since there clearly is free energy to be released. In fact, although the motion is caused by an electromagnetic wave, the resulting flow actually appears to be unstable to purely electrostatic modes, with growth rates much faster than the laser period, as we will demonstrate presently.

To give a somewhat quantitative analysis of the instability, we follow the method of Mikhailovskii [82]. The equations we will need are the continuity equation

$$\frac{\partial n}{\partial t} + \nabla \cdot (n\vec{V}) = 0, \quad (2.6)$$

Euler’s equation for a plasma (ignoring the magnetic field)

$$\frac{\partial \vec{V}}{\partial t} + (\vec{V} \cdot \nabla) \vec{V} = \frac{e}{m_e} \vec{E}, \quad (2.7)$$

and the Poisson’s equation for electrostatics

$$\nabla^2 \phi = 4\pi e n_T \quad (2.8)$$

²⁰ We adopt here the habit, which seems common in the literature, of describing plasma instabilities as ‘Kevin-Helmholtz like,’ ‘Rayleigh-Taylor like,’ et cetera. However we wish to emphasize that the simile is little more than nominal; strictly speaking those instabilities are well defined, and usually involve phenomenology which is often aren’t pertinent to laser plasmas (for instance viscous fluids.) More specifically, the analogy is usually limited to the geometry of the physical situation (Rayleigh-Taylor meaning longitudinal pressure gradient, and Kevin-Helmholtz being a transverse velocity gradient, for instance) and as such says nothing about what allows the instability to grow. None-the-less we find the phrasing is useful for forming a mental image of the problem, so long as it is remembered that it’s really just a mnemonic.

where in the first two equations n and $\vec{V}$ are understood to be the parameters for each species, and in the last equation n_T is the total density for all species. The electric field is related to the potential by $\vec{E} = \nabla\phi$.

Before $t = 0$ we take the plasma parameters and electric field to be in some initial state (which may depend on time), $n = n_0$, $\vec{V} = \vec{V}_0$, and $\vec{E} = \vec{E}_0$, which by themselves satisfy our two equations and so if left undisturbed are stable. However at $t = 0$ we do disturb them, so that their new parameters are given by $n = n_0 + n_1$, $\vec{V} = \vec{V}_0 + \vec{V}_1$, and $\vec{E} = \vec{E}_0 + \vec{E}_1$. Substituting these values into equations 2.6 and 2.7, dropping all terms which are the product of more than one perturbed value, and using the the same equations with initial equilibrium values to eliminate the zeroth order terms, we arrive at

$$\frac{\partial n_1}{\partial t} + \nabla \cdot (n_1 \vec{V}_0 + n_0 \vec{V}_1) = 0 \quad (2.9)$$

$$\frac{\partial \vec{V}_1}{\partial t} + (\vec{V}_0 \cdot \nabla) \vec{V}_1 + (\vec{V}_1 \cdot \nabla) \vec{V}_0 = \frac{e}{m_e} \vec{E}_1 \quad (2.10)$$

To apply this to a laser reflecting at a surface, we assume the plasma to have some velocity V_0 in the z -direction (and none in the other directions), which has a dependence on the x -direction such that $\partial/\partial x \not\rightarrow 0$ even for the zeroth order or equilibrium conditions. (Note the laser electric field does not appear in these equations and so does not enter directly into their solution.) The equilibrium plasma density may or may not vary in x as well, but all zeroth order quantities have translational invariance in y & z . With these values for the equilibrium parameters, equations 2.9 & 2.10 reduce to

$$\frac{\partial n_1}{\partial t} + n_1 \frac{\partial V_0}{\partial z} + n_0 \frac{\partial V_{1y}}{\partial y} + n_0 \frac{\partial V_{1z}}{\partial z} + \frac{\partial(n_0 V_{1x})}{\partial x} = 0 \quad (2.11)$$

$$\frac{\partial \vec{V}_1}{\partial t} + V_0 \frac{\partial \vec{V}_1}{\partial z} + V_{1x} \frac{\partial V_0}{\partial x} \hat{z} = \frac{e}{m} \nabla \phi \quad (2.12)$$

At this point these equations are generally valid. To attempt a solution however we will need to make some assumptions. First we assume that the zeroth order quantities are only slowly varying in time; we can then take time derivatives of the perturbed quantities to be harmonic and make the substitution $\partial/\partial t \rightarrow -i\omega$. This will be valid in our case provided the growth rate is sufficiently greater than the laser frequency. We take the dependence in the y -

and z-directions to be sinusoidal as well, so that we can make the substitutions $\partial/\partial y \rightarrow ik_y$, and $\partial/\partial z \rightarrow ik_z$.

We now arrive at (expressing the vector equation as three separate equations):

$$0 = -i(\omega - k_z V_0)n_1 + ik_y V_{1y}n_0 + ik_z V_{1z}n_0 + \frac{\partial}{\partial x}(n_0 V_{1x}) \quad (2.13)$$

$$-\frac{e}{m} \frac{\partial \phi}{\partial x} = -i(\omega - k_z V_0)V_{1x} \quad (2.14)$$

$$-i\frac{e}{m}k_y\phi = -i(\omega - k_z V_0)V_{1y} \quad (2.15)$$

$$-i\frac{e}{m}k_z\phi = -i(\omega - k_z V_0)V_{1z} + V_{1x}\frac{\partial}{\partial x}V_0 \quad (2.16)$$

If we now use these equations to solve for the perturbation to the density n_1 , and insert this result into equation 2.8, we arrive at

$$\frac{\partial}{\partial x}(\epsilon_0 \frac{\partial \phi}{\partial x}) - (k_y^2 + k_z^2)\epsilon_0\phi = 0 \quad (2.17)$$

where $\epsilon_0 = 1 - \omega_p^2/(\omega - k_z V_0)^2$, and again V_0 is a function of x , as may be ω_p .

To determine a growth rate we must be more specific about the behavior of the zeroth order plasma behavior. We consider the skin depth to be broken into two regions with constant but different V_0 's and with the same density. Each region can be imagined to have a thickness of $\sim c/2\omega_p$ (though this isn't necessary for the derivation). The velocity in the outer region is given by $V_{01} = eE_0/m_e\omega_0$ where E_0 is the peak electric field at the surface and is related to the laser vector potential by $eE_0/m_e c\omega_0 = 2a_0\omega_0/\omega_p$. The velocity in the inner region we take to be half this. Following the analysis of [82] this leads to a dispersion relation of

$$0 = 2 - \frac{\omega_p^2}{(\omega - kV_{01})^2} - \frac{\omega_p^2}{(\omega - kV_{02})^2} \quad (2.18)$$

which predicts an instability with a maximum growth rate at a wavenumber of

$$k_z = \sqrt{3/2}\omega_p/V \quad (2.19)$$

with a growth rate of

$$\gamma_{max} = \omega_p/2\sqrt{2} \quad (2.20)$$

where $V = V_{01} - V_{02}$.




FIGURE 2.18: Phase spaces of forward (a) and transverse (b) momentum vs. forward position for a plasma interacting with a laser but with no magnetic force on the particles. In the absence of any instabilities, the motion of the particles would be fluid and in the transverse direction only and no heating would occur. That heating in both the transverse and forward direction does occur shows that the flow is indeed unstable.

To verify that the real flow is in fact unstable with a high enough growth rate to explain the heating we see in our two-dimensional runs, we ran simulations with similar parameters (though with fixed ions) but where we have modified OSIRIS so that the particles feel no magnetic field. The fields are still updated correctly from the the full set of Maxwell's equations and the particle quantities, so the laser will still be absorbed and reflected with the skin-depth. But by zeroing out the magnetic force we will remove the non-linear movement of the surface which was not include in the theory just described, and prevent the plasma from being unstable to anything other than electrostatic-like modes. We initialize the plasma with a temperature of 1° K , so we are certainly in the cold regime²¹.

Phase space results are shown in figure 2.18 after roughly two cycles of laser illumination (images are taken approximately at the time of a node in the laser field.) As can be observed, strong transverse heating has already taken place, with normalized momenta reaching above unity in both directions. Heating in the forward direction is also apparent (2.18.a); although the peak velocity is somewhat less than for the transverse direction (2.18.b) – and as we have been emphasizing longitudinal heating plays a minimal role in the acceleration mechanism.

²¹ This is of course unphysically cold and it may be unclear why any temperature at all is used; the reason is that otherwise there is no disorder to seed the instability

This is noteworthy because in the absence of the electrostatic instability the momentum in the forward direction would remain precisely zero without the magnetic force. Measuring the temperature at the surface a value of roughly 75 keV is found. It seems clear therefore that this “slipping streaming” instability is able to break the translational invariance in x_2 sufficiently to allow retention of laser energy between cycles and by doing so to provide the bootstrap heating necessary to access the forward particle acceleration at the surface.

The single species electrostatic picture described so far seems to explain well the early time heating we see in full PIC simulations. However after a period of a few 10’s of laser cycles, we begin to see a divergence between fixed and mobile ions runs. Specifically, up to this point the spectrum of accelerated electrons (along with other diagnostics) appear equivalent between the two cases; but afterward the spectrum seems to peak and level off for fixed ions. A likely reason for this is that acceleration mechanism, by nature of the fact that it selectively pushes the highest energy particles into the bulk, will have a cooling effect on the plasma surface; eventually the heating and this cooling are in balance, and a steady-state is reached.

In contrast, for mobile-ions the surface continues to heat and the peak particle momentum continues to grow, up to $2a_0m_e c$. Coincidentally with this we see a ‘pocketing’ or ‘rippling’ of the plasma surface, as can be seen in figure 2.19.a. Recently Wan et al. [83] attributed this to an instability two-stream like instability (OSTI) in the skin-depth layer. We note that it is quite possible that the “slipping stream” instability could seed the OSTI. The ion instability allows perturbations in the electron density to persist through laser cycles. The resulting magnetic field filaments caused by the forward current of the accelerated electrons and the neutralizing return current, though still weak at this time, may also play a part. As we stated earlier, many mechanisms are likely at play at the same time throughout the interactions. Developing a more exact analysis of the slipping stream instability and how it couples to the OSTI is an area for future work.

We believe these surface perturbations lead to an increase in heating and acceleration via two mechanisms. The first, which is related to heating, can be thought of as being similar to Brunel heating, in that the laser is no longer normal to the surface at all points. How-




FIGURE 2.19: The plasma density at the surface (a) after $60fs$ of laser illumination, and the magnetic field (b) at the same time (note the spatial scales are not the same). The surface remains very nearly smooth, but structures at roughly the scale of the skin depth are beginning to appear. Even though variations this small leave the plasma essentially optically flat, the standing wave structure at the surface is slightly modified, leading to local reductions in the magnetic field which may allow some electrons to more easily enter the laser fields.

ever, in this case it isn't that the electrons are being disruptively pushed into the plasma, as Brunel described, but rather that the electrons which are traveling parallel to the (average) plasma surface see a strongly varying density in time and space, breaking the translational invariance and harmonic motion of their orbits and allowing them to retain energy after each cycle, i.e. heat. G.E. Kemp et al. [76] observed a similar effect when they initialized a plasma with a small scale ($\sim c/\omega_p$) surface irregularity, i.e. a large increase in early time absorption compared with a perfectly flat plasma was seen. They found this surprising because a perturbation of this scale is by most measures optically flat. However the point here is that what is pertinent in this case is not what the laser 'sees,' but rather what the electrons do; and for them the important scale length is not the laser wavelength but rather the skin depth. By that metric the surface bumps are indeed significant, and so it is no longer surprising that they have an effect on absorption.

The second effect is that, even though surface remains optically smooth to a first order approximation, still the field very close to the plasma can become modified by the density fluctuations, with the scale of the field modifications $\sim c/\omega_0$ even for fixed ions. Such a field distortion can be seen in figure 2.19.b. What this means is that certain sections of the surface may see a lower peak magnetic field than in the perfect standing wave, allowing electrons to more easily escape into the accelerating electric field and so gain energy. The change in the standing wave field structure is subtle and hard to detect, and the consequent effect on the electrons is likely a modest one, but it is still imaginable that it plays a part in the enhanced acceleration; it is, however, more important later in time, when the surface is more grossly perturbed.

2.2.5 Circular Polarization, and s-Polarization

The absorption physics is quite different, as compared to p-polarization, for either circular polarization, which has real and useful applications; or so-called s-polarization, which doesn't actually exist.

We remind the readers that s-polarization has a much different meaning in simulations




FIGURE 2.20: *Forward-momentum and forward-position phase space (a) is shown for a two-dimensional plasma after 500fs interacting with a laser with its electric field out of the plane of the simulation (‘s-polarization’). Exactly zero particles are accelerated to relativistic energies. The surface of the plasma at the same time is also displayed (b), which shows that the plasma remains perfectly flat.*

than it does in reality: an s-polarized normally incident wave is one in which the wave’s electric field is pointed in the unphysical, or ‘ignorable,’ direction. This of course is a purely artificial geometry, one which has no analogy to a three-dimensional scenario (or a one-dimensional one, for that matter.) As such, any difference between s- and p-polarization is non-physical and hence not particularly interesting by itself. However, such comparisons permit clear elucidation of important physics, and so it is worth looking at them briefly.

The results are summarized in figure 2.20.a; in short, there is no absorption for s-polarization. The reason for this is simple. With the magnetic field now in the plane of the simulation, the particles must have thermal velocities out of the plane in order to use that field to rotate into the vacuum region where the accelerating electric field is; but heating is impossible out of the plane because transverse canonical momentum remains rigorously conserved in this direction. Absorption physics with s-polarization is essentially much the same as in a 1D geometry. Interestingly, heating also does not happen to an appreciable amount even in the plane of the simulation. This makes sense as the instability just de-

scribed in section 2.2.4 depended on having a k -vector in the direction of fluid flow; indeed there seems to be no surface instability in this geometry, as the surface stays almost perfectly flat even on picosecond timescales - as can be seen, at least initially, in figure 2.20.b.

Circularly polarized laser light is also unable to accelerate electrons (so long as the surface remains flat), but for quite different reasons. This is despite the fact that surface heating still occurs because there is still a component of the electric field that is in the plane of the simulation. This heating can be seen in figure 2.21.a which shows the electron p_1 vs. x_1 phase space for a plasma under circularly polarized laser illumination. However, as can be seen, despite this heating no electrons are able to be accelerated to MeV energies as they can under linear p-polarized illumination. This can be quickly understood by recalling the acceleration mechanism. In the standing wave structure of incident and reflected circularly polarized light, the surface magnetic field rotates in time but its magnitude is constant. No matter when a particle reaches the surface, it will see a strong reflecting magnetic field. As a result there is no phase where an electron can ride a decreasing magnetic field out to where the electric field is at a maximum, as is the case for a linearly polarized laser.

The lack of significant electron acceleration and the resulting lack of bulk plasma heating is essential for the so-called ÒRadiation Pressure AccelerationÓ (RPA) scheme for generating mono-energetic ion beams. Such a beam is shown in fig. 2.21.b. In RPA an intense laser impinges on a plasma. The ponderomotive force pushes the plasma electrons forward thereby generating an electric field that pulls the ions with them. The laser acts like a piston which then reflects the ions to twice its inward velocity. This velocity was derived earlier and is based on the momentum balance between the laser and the plasma ions. This effect is similar to the interaction we saw in 1D simulations (with linear polarization) with cold electrons and mobile ions, c.f. 2.8, only in that case it was an artifact of using 1D simulations, while here it is a physical effect.

If the surface electrons are accelerated by the 2a0 mechanism, two processes can occur that prevents the RPA process from occurring. When the electrons are continually accelerated forwarded with relativistic energies then the bulk plasma and not just the surface is heated. When this occurs the ion acoustic velocity becomes comparable to the inward velocity of




FIGURE 2.21: The effects of circular polarization on laser absorption are shown for two different simulations. In frame (a), the forward-position versus forward-momentum electron phase space is shown for a plasma after $250fs$ of $a_0 = 6$ laser illumination; similar to s -polarization in two-dimensions, no acceleration and little heating is visible. Frame (b) shows again $p_1 x_1$ phase space, but for the ions (deuterium) rather than electrons, after $30fs$ of $a_0 = 100$ laser illumination. Were the electrons able to gain energy from the laser, they would launch an ion-acoustic shock which would interact with the ions within the plasma; here instead the ions are accelerated directly at the surface and so are monoenergetic (with a velocity twice the compression velocity). This is known as the Radiation Pressure Acceleration mechanism.

the piston and an ion acoustic shock can be launched. In this case, Silva et al. [84], showed that the acoustic shock can trap and reflect a much smaller number of ions. In addition, the acceleration of the electrons increases the laser absorption thereby decreasing momentum being imparted to the plasma ions. The results in this dissertation shed light for the first time as to why RPA only works for circularly polarized light.

2.3 Absorption in the Hydrodynamic Regime - Picosecond Timescales and Micron-Scale Preplasmas

In the previous section, we studied the $2a_0$ laser absorption mechanism present in self-consistently heated two-dimensional plasmas. However although - as we have made pains to stress - the heating mechanism is inherently multi-dimensional, we have restricted our description and examples to situations where the bulk plasma density (or, synonymously due to quasi-neutrality, the ion density) is initially transversely smooth and unperturbed. In addition, the 1D density function was restricted to one of the simplest possible, that is a step function. It might be accurate and illustrative to describe the previous geometry as one having two-dimensional electrons, but one-dimensional ions. We refer to this as a quasi-1D geometry.

For some experiments and targets, the description given so far may in fact be close to a complete picture. Specifically, if the experiment has some combination of a short²² laser pulse, a high laser contrast ratio, an initially smooth (and therefore, almost certainly, cold) target surface, and a high target mass density, the interaction may start and end in a regime well described by such a quasi-1D geometry.

However, there are of course many other scenarios that involve the absorption of high-intensity lasers which relax any subset of the previous restrictions. Perhaps most pertinently, the laser pulses proposed in Fast Ignition must have pulse lengths on the order of picoseconds, which will permit the ions to evolve hydrodynamically; in addition in the initial FI concept the ‘target’ would be a hydrodynamically compressed hydrogen plasma, which one would be hard pressed to imagine as presenting a perfectly smooth critical surface. The mechanism described here is probably more relevant to the embedded gold cone model of FI than to the symmetric target model. The natural question then becomes, does the mechanism so far described have relevance to the more complex picture, where ion motion including target

²² ‘Short’ in this context however being a parameter window rather than a limiting case, ‘short but not too short’ that is - for femtosecond or shorter pulses, this model is not expected to be relevant (as heating does not happen on that timescale). See for instance the work of Baeva et al.[85, 86] for a mechanism pertinent to that regime.

expansion is included? Or is it limited to a small subset of experimental conditions?

We argue – and will demonstrate presently – that although the physics of laser absorption is indeed more complicated when the ions evolve on hydrodynamic time scales, the $2a_0$ mechanism described here remains in most cases the dominant mechanism, especially if one is most interested in the generation of moderately energetic (a few to 10 $\tilde{O}$ s of MeV) as is needed in fast ignition. We will show that it is just one other process becomes very pertinent and that it generates very energetic electrons (with relativistic gamma factors greatly exceeding $2a_0$). This process occurs when the plasma profile acquires a low density shelf. Such a shelf – as well as a critical surface which rises sharply to high density but which has a ‘rippled’ shape – arises when an intense laser compresses a plasma with a less sharp density gradient, for instance one formed from a solid density material being ablated by a lower intensity laser; or when an intense longer pulse laser impinges on an initially sharp interface.

2.3.1 Hydrodynamic Evolution of an Initially Flat Surface

We begin by examining a simulation with the following parameters. As before, all normalized parameters will be scaled the laser frequency, ω_0 - which we will refer to as having a wavelength of $1\mu m$ for simplicity - along with the physical constants (m_e , e , and c .) We run OSIRIS in a 2 $^{1/2}$ D slab geometry, with square cells of width $0.05c/\omega_0$. The time step is just shorter than required by the Courant condition, at $\delta t = 1/57 \simeq 0.0175$. The simulation box is $3\mu m$ wide and periodic in the transverse direction, and $250\mu m$ long in the longitudinal direction; the simulation grid is therefore 63000×756 cells. In the rear $10\mu m$ of the box we have altered physics as we will describe briefly. All particles are initialized with the same weight, which corresponds to $103n_0$ when there are four probabilistically and ions per cell. The left-most $160\mu m$ of the box is initially a vacuum region, and the plasma density rises rapidly from 0 to $100n_c$ within a cell. We use a sinusoidal variation on the average number of particles per cell, so that the density varies from $93n_c$ to $103n_c$ in the plasma region, for reasons described earlier. The plasma is $80\mu m$ deep, with another $10\mu m$ region mentioned

earlier. Ions are initialized at the exact same locations and with the exact same charge as electrons, and have a charge-to-mass ratio of 3643.8 (deuterium), and both ions and electrons are initiated with a temperature of 1 keV. The laser is launched from the left wall with a normalized amplitude of $a_0 = 6$, the k -vector points in the x_1 direction, and it is p-polarized so that the laser electric field is in the plane of the simulation.

Throughout the simulation box, we use a particle splitting method on the high energy electrons to reduce the effect of macro-particle stopping (as described in Chapter 4, with up to five binary splits for a maximum effective particle-per-cell count of 32 times the initial count (that is, 128 particles-per-cell); specifically, the particles will be split each time their relativistic factor γ exceeds these values: 1.4, 1.8, 2.2, 2.6, 3.0. At each particle splitting event, we add a small amount of transverse momentum to each particle so that they no longer overlap in phase space. Here we use a random number for the momentum such that the magnitude is 1% of the initial momentum in the lab frame. Momentum is conserved exactly but a small amount of heating does occur²³.

The goal of these simulations is to study the absorption of the laser over time scales for which the ion density profile evolves. Over such time scales significant numbers of accelerated electrons penetrate into the plasma. In addition, for fast ignition, the target is very large and it is not possible to model the entire target even in two dimensions. Therefore, if the boundary conditions in the simulations are not properly addressed then the plasma conditions where the laser is being absorbed can be significantly altered.

A major issue is the reflux of the accelerated electrons. When only simulating a small region of a large target then large numbers of accelerated electrons will hit the simulation boundary. These electrons can then reflux leading to the entire simulation volume heating up to unrealistic temperatures. We note that for small targets refluxing and the heating of the plasma is a real effect. In order to prevent refluxing we have implemented a “artificial” damping region in the right most $10\mu m$ of the box. We compare each particle’s kinetic energy to the local temperature to determine if the particle is “hot”, i.e., above a threshold value.

²³ Note that it is actually a very small amount: less than one part in ten thousand at each event on average.

This threshold energy is called the damping temperature. We then “cool” all particles that are above the damping temperature with a probability which provides a specified stopping distance. More details are given in Appendix B.

In addition to the particle damping algorithm, we use a particle merging algorithm for thermal particles in the same $10\mu m$ region. This is done in order to reverse the splitting which occurred when the particles were accelerated. If we do not use such an algorithm, the background plasma in the damping region becomes over-populated with smaller particles, and that region can become the computationally dominant part of the simulation, empirically slowing the simulation down by a factor of 8 or more. The algorithm uses the same list of kinetic energies as the splitting algorithm to determine when to merge, but it merges the particles at a lower energy than that at which they were split (specifically a factor of 4 less) to prevent particles which are near the threshold from merging and splitting multiple times. Only particles close in both position and momentum space are allowed to combine. Although it is possible to conserve both energy and momentum by combining 3 or more particles [87], we elect to use a simpler two-particle merge which conserves momentum only, since our region is already (purposefully) non-energy conserving. The routine also spatially balances the sites of particle merging, by keeping track of previous merge events and preventing more from happening in regions in which a disproportionate number of events have happened already. This is done primarily to prevent merging from happening predominantly on the edges of the merge region.

Figure 2.22 shows the electron spectrum and figure 2.23 the plasma density roughly 300 femtoseconds after the initial laser illumination. This time period is interesting because it shows the transition between the initial, simple geometry, and the somewhat more complicated absorption physics which happens later in time. Up to about 250 femtoseconds, the surface remains optically smooth, and the absorption is from the 1D physics for a hot plasma, as can be seen by the fact that in fig. 2.22.a the maximum energy for the electrons for $x_1 > 1200c/\omega_0$ is $\sim 2_a0 = 12$. However, as can be seen in figure 2.23, by $t\omega_0 = 1600$ the surface has developed laser-wavelength scale structures, specifically ocean-wave-like asymmetric hooks, terminating in sharp points. These hooks have two consequences. One, their chirality




FIGURE 2.22: *Electron spectrum 300fs after laser illumination. Fig. 2.22.a shows the whole plasma region; the higher maximum energy of the particles accelerated later (and so less deep into the plasma) can be clearly seen. Fig. 2.22.b shows a zoom around the surface, displaying the same $2\omega_0$ jets as earlier in time, though now those from each laser half period has distinct shape.*




FIGURE 2.23: *Plasma density at the surface 300fs after laser illumination. The previously smooth surface has developed bumps and hooks, an underdense plasma can be seen in front of the target, and a high density shock-like structure is traveling into the plasma.*

breaks the symmetry of the laser-plasma interaction, which results in different absorption in each of the two half-cycles of the laser. This can be clearly seen in figure 2.22.b, where the half-cycle beam pulses are alternately strong and weak. This handedness arises from local spontaneous symmetry breaking, and although their absolute orientation is of course random and also likely to change over time (that is on a time scale much longer than the laser period) they can lead to 1ω harmonics whereas the perfectly smooth surface would only give 2ω . The surface structure also complicates the laser reflection so that it is no longer a planar standing wave, allowing for peak acceleration of greater than twice the quiver velocity under favorable conditions, as can somewhat be seen in figure 2.22.b - for instance the light-blue peak at $x_1 = 1015$ reaches roughly $15 m_e c^2$.

However the highest energies seen in the spectrum are more than double those from the smooth surface, with a low density group extending up to gammas of 30 or greater. Such energies are far too high to be explained by a local focusing of the field. Rather this is a signature of a second acceleration mechanism, one which is related to another feature visible

in figure 2.23 - the presence of an underdense plasma in what used to be the vacuum region in front of the target (roughly $1000c/\omega_0 < x_1 < 1010c/\omega_0$).

We examine the source of these electrons using particle tracks, which can be seen in figures 2.24 to 2.29 . The first figure shows the particles we will be tracking, and demonstrates the selection rules we have chosen. The particle must be near but not in front of the plasma surface at $t\omega_0 = 1600$ (i.e. $x_1 > 1010c/\omega_0$), and that the normalized forward momentum be greater than 20. Based on these rules we selected 150 particles. The following figures (2.25 -2.29) show the tracks of the selected particles at snapshots separated by $t\omega_0 = 15$ preceding the time of selection, superimposed on the plasma density. These tracks then help to identify the acceleration mechanism for the “super hot” electrons. The position of the particles at the time of each frame are indicated with yellow dots, the time history of their motion is shown with a tail behind each dot, which is also color-coded to show the particle energy at the previous times and positions.

We would like to emphasize that the selection criteria illustrated– $p_1 > 20m_e c$ and $1010c/\omega_0 < x_1 < 1040c/\omega_0$ – are the only restrictions we place on the particles. Otherwise we pick the electrons to track completely at random, and all 150 selected tracked particles are plotted in the figures that follow.

The first point to notice upon examining the entire data set for the tracks is that every single electron spends time in the underdense region in front of the plasma (the area which started out as vacuum). Some have already exited the plasma by the first frame (which is taken 30 fs before the “final” time), and the last of the 150 are exiting the plasma by the third frame, just as the first group of electrons begin to re-enter the bulk plasma. This behavior should be contrasted with the movement we saw earlier in the particles which were accelerated by the $2a_0$ mechanism at a flat surface. In the $2a_0$ process the particles exited the plasma as well, but they did so in a smooth manner in phase with the laser and re-entered the plasma right away. They only penetrate $\sim \pi c/2\omega_0 = \lambda_0/4$ into the vacuum where the E field in the standing wave is at a maximum. Each accelerated electron does this in phase with the wave. Here instead the particles are able to leave the plasma entirely and enter the previous vacuum region, and their motion is not smooth or in phase. Rather the particles

execute chaotic motion both in position and momentum space. These particles decouple from the dense plasma to become part, at least for a time greater than the laser period, of the underdense plasma shelf which exists in front of the original surface.

Once we have identified that all the particles execute motion inside the underdense plasma shelf, the next interesting issue to investigate is what determines and how the particles exit the plasma and enter the laser field region. They do not exit the plasma in a spatially uniform manner; rather virtually every electron exits the plasma by traveling along the density ‘spikes’ pointing out into the vacuum region²⁴. In fact the majority of the particles exit through just one of these spikes (the second from the left in these frames in which there are four spikes). Again this is in contrast with the behavior of tracks for the smooth surface “1D” like behavior. In that case all points along the surface were equivalent since the fields were translationally invariant in the transverse direction. As a result, all particles see the anti-node of the magnetic field of the standing wave which prevents electrons from entering the vacuum (except for a few who penetrate $\lambda_0/4$ but return immediately). When the ions move the smoothness of the interface, and hence the fields, is broken which breaks the translational invariance. The plasma and hence the fields become localized into regions λ_0 wide. The fields are focused into the low density regions. Therefore, the particles escape through the density spikes where the fields (especially the B field) are low. In essence, the spikes then extend into the vacuum region and are the source of the underdense shelf²⁵.

We can also understand through the tracks the way the particles move in the fields and how they re-enter the plasma. In the standing wave, the particles move in a stochastic and overlapping fashion, mostly in transverse paths but sometimes longitudinally as well,

²⁴ If one looks carefully there is one exception to this rule visible in these particle tracks, a dark purple track which can be seen exiting the plasma in the second frame between the left and middle spike. However though not shown here this is the exception which proves the rule: in fact this particle did leave along one of the spikes, but at an earlier time. It then reentered the plasma quite energetically but with grazing incidence, and subsequently used that energy to exit the plasma again (as seen here.)

²⁵ Sometimes it is stated that an underdense shelf is created by the diffusive movement of the ions into the vacuum region. This is not a meaningful way to explain the evolution of the plasma; the ions do not feel the longitudinal force of the laser field, but rather only respond to the electrostatic fields setup by the movement of the electrons. Once the electrons are able to enter the vacuum region, the ions then follow them to maintain quasi-neutrality (again, doing so without difficulty, because they are not affected by the laser); but their movement is a consequence of the escape of the electrons, not a cause.

both gaining and losing energy throughout their motion. Eventually the particles hit the surface at various locations²⁶ and angles, after which they travel ballistically into the plasma. Although all the electrons appear to re-enter the plasma around the same time, this is just a consequence of the original selection rules; more illuminating is how long the electrons spend in the vacuum, which is a much broader distribution.

2.3.1.1 Acceleration mechanism for electrons of unbounded energy

For laser pulses with longer laser pulses we have identified two groups of energetic (accelerated) electrons, a hot and a super-hot group. The hot group has energies approaching $2a_0m_e c^2$ and they are generated through the 1D mechanism discussed earlier. However, as we have just described for hydrodynamic times scales the plasma surface ripples and the laser is focused into the low density regions. The plasma then leaks into the vacuum through the high density spikes thereby creating an underdense shelf. The $2a_0$ mechanism at the surface is slightly modified due to the density modulation, but it still occurs within $\lambda_0/4$ of the surface. However, the underdense shelf is the source of the high energy, i.e., super-hot electrons.

The acceleration mechanism for the super-hot electrons is stochastic acceleration of electrons in the standing wave pattern. The electrons quickly gain and lose many multiples of the quiver velocity (momentum) without any particular limit. Eventually, however, their stochastic path intersects the surface of the plasma and they then travel along straight line trajectories into the bulk plasma.

The stochastic acceleration of electrons in a combination of forward and backward propagating electromagnetic waves has been investigated in the past [88, 89, 90, 91, 92]. Stochastic acceleration occurs when particles interact with two waves with different phase velocities. For the purposes of this work, we note that when an electron interacts with a single light wave in vacuum it can gain an energy which scales as a_0^2 . Its energy depends on its initial conditions.

²⁶ Although less so than for exiting, the electrons also appear to have some preference to enter the plasma along the spikes as well; this may be a simple geometric feature, it being more likely for a ray to strike a hill than a valley.




FIGURE 2.24: *Particle tracks for high momentum particles.*

Above: The particle selection. The window parameters show the selection rules, the black dots the selected particles.

On the following pages: The tracks for the selected particles, superimposed on the plasma density, in five frames roughly $7.5fs$ apart. Yellow dots indicate the location of the particle at the current time; the tails show the location over the previous $10/\omega_0$, and are color-coded by kinetic energy.




FIGURE 2.25: Tracks at $t\omega_0 = 1540$




FIGURE 2.26: Tracks at $t\omega_0 = 1555$




FIGURE 2.27: Tracks at $t\omega_0 = 1570$




FIGURE 2.28: Tracks at $t\omega_0 = 1585$




FIGURE 2.29: *Tracks at $t\omega_0 = 1600$*

The counter propagating light wave can also interact with the electron which effectively can bootstrap its energy by changing the effective initial conditions for the electron when it is interacting with the forward moving light wave. Therefore, electrons which can interact with a traveling wave in a region with many laser wavelengths long can gain significant energy. In the present case, the situation is even more complex, as due to the uneven structure of the plasma surface, the reflected wave can be much more uneven than the incoming wave, and so the resulting motion of the electrons more chaotic. With respect to fast ignition, the stochastic process generates electrons significantly higher than those that can be stopped in the fuel. Therefore, this process is undesirable. We believe that a careful analysis of the motion of electrons in a traveling wave and of those in the underdense shelf still needed.

2.3.1.2 Plasma expansion into vacuum region

The spectrum of accelerated electrons, then, is in the simplest sense a direct effect of the longitudinal profile of the plasma. The ‘modestly’ energetic electrons, i.e. ‘hot’, those up to $2a_0m_e c$ (and hence those most pertinent for Fast Ignition, among other uses) are created by the interaction of the laser with a sharp, steep interface on a highly over-critical plasma, as we saw earlier. By contrast the most energetic, super-hot electrons, which represent ‘wasted’ energy in the FI concept, are caused by any underdense plasma which may be present. It is useful, therefore, to look at how the plasma profile evolves over the length of what might be a typical laser pulse, which is of order a few picoseconds. Studying the evolution of the surface and plasma shelf will also inform the validity of ignoring the effects of any laser prepulse or compressional scale length, as we will return to shortly.

In figure 2.30 we show the plasma density, averaged transversely, on a logarithmic scale, from the time of the initial laser illumination out to approximately 2.6 ps. There is of course also structure in the transverse direction, especially around the surface. However, this diagnostic will provide information on how the plasma surface moves inward and how the shelf forms. Figure 2.30.a, which includes the entire initial vacuum region of the simulation, shows the expansion of the underdense plasma shelf. The initial source of electrons for this




FIGURE 2.30: The density of the plasma averaged transversely as a function of time. a) The expansion of the plasma as an underdense shelf (green and red regions) in front of the compressing overdense plasma (magenta) can be seen. b) A zoom around the surface shows the preservation of a sharp transition from underdense (green) to overdense (red fading to yellow) regions.

plasma is again the micron-scale (laser wavelength scale) spikes on the plasma surface which ‘pierce’ the surface field and allow the electrons to leak out. Once in the underdense region the electrons no longer feel the laser pressure, as the plasma is transparent for the laser, and so they expand out ion-acoustically. The fluid velocity of the ions is a linear function of the distance from the plasma surface (not shown here), and so the underdense gas expands exponentially in time, as can be seen. The density of this shelf therefore remains well below even the non-relativistic transparency density, which is on order of one-tenth the critical density or only a few percent of the relativistic critical density).

Figure 2.30.b, shows the same data, but now zoomed in around the high density surface region. The most important point is that, although the critical surface moves constantly to the right, due to a combination of the compression inward from the radiation pressure and etching outward, the transition from underdense to highly overdense remains sharp. The color axis here is chosen so that there is a quick transition from green to red at the non-relativistic critical density. In addition, although faint, a shockwave is also seen to be launched into the plasma as the separation between the light and slightly darker yellow region. The density behind the shock is about $150n_c$ (i.e. 1.5 times the initial density).

Figure 2.31 shows 2D snapshots of the density profile at the surface at four sequential times. As we expected from what we saw in fig. 2.30.b, here we see that although the surface is constantly moving to the right, it maintains the same sharp transition from low to high density throughout. Indeed, as could be anticipated, the sharpness of the interface is even more distinct in these 2D pictures, as they aren’t smoothed over by the transverse averaging. Furthermore, the general two-dimensional shape of the surface also remains constant throughout the simulation. On scales longer than the laser wavelength the surface stays relatively flat, with angles of only a few degrees from normal, while on scales similar to the laser wavelength it stays ‘spiky’ throughout. This is consistent, respectively, with the continued ability of the laser to generate energetic electrons through the $2a_0$ mechanism - as we’ll see is true shortly - and the constant expansion of the underdense region in the vacuum, as seen in fig. 2.30.a.




FIGURE 2.31: *The critical plasma surface shown at four sequential times, separated by approximately 650 femtoseconds. Although the surface moves constantly to the right, due to a combination of compression and etching, the basic characteristics of the plasma - a sharp transition from underdense to highly overdense, shallow angles above micron scale-lengths but possibly with sharp spikes at sub-micron scales - remain the same throughout.*

2.3.1.3 Heat flux spectrum of accelerated electrons

Examining in detail the spectrum of accelerated electrons can reveal the interplay between the two different absorption mechanisms. Figure 2.32 shows the total forward heat flux carried by electrons as a function of their relativistic gamma factor, as a function of time (taken as an average both transversely across the box and longitudinally in a volume 6 microns behind the original surface and 18 microns thick.) On this multi-picosecond timescale, the period where only the smooth-surface acceleration mechanism is present appears as only a brief transitory event lasting a few hundred femtoseconds at the beginning of the interaction (until $t \sim 1500\omega_0^{-1}$, which is $\sim 500\omega_0^{-1}$ after the laser reaches the surface). Immediately after this transition time, electrons begin appearing with momenta above $2a_0m_e c^2$, and with energies even beyond the scale shown here. These super-hot electrons are correlated with the appearance of an underdense region. These electrons are generated through the interaction in the standing wave.

However, the lower part of the spectrum (the hot electrons) which is caused by the $2a_0$ absorption process are still visible throughout the simulation (the white-purple-red area of the plot). And, more critically, the majority of the heat flux continues to be carried by electrons with $\gamma < 2a_0$, demonstrating that the energy absorption is still primarily through the $2a_0$ mechanism. (A burst of higher-gamma absorption can be seen to have happened briefly around $t = 2000/\omega_0$; this corresponds to the presence of a unusually dense area of pre-plasma, which can also be seen in fig.s 2.30.b and 2.31.a.)

To illustrate more clearly the presence and relative importance of the hot and super-hot electrons we plot in figure 2.33 snapshots of the different moments of the electron distribution function, on an energy against longitudinal position plot. We use $\log_{10}(\gamma - 1)$ for the y-axis and x_1 for the x-axis. In fig. 2.33.a, .b, & .c we show the 0^{th} moment, i.e. plasma density, the 1^{-n} moment, i.e. current density, and the 3^{rd} moment, i.e. the heat flux, respectively. Each quantity is averaged transversely at time $t = 3600\omega_0^{-1}$, which corresponds to roughly the middle of figure 2.31. Fig. 2.33.a shows that the bulk plasma is at a temperature of 5-50 keV; the relativistic electron density isn't even visible with this linear color map. In fig. 2.33.b we




FIGURE 2.32: Time series of the heat flux spectrum, that is the total forward heat flux carried by electrons as a function of their γ . Although energies above twice the quiver velocity are only strictly excluded for the brief period at the beginning of the graph (corresponding to the time when the surface remains smooth), the heat flux is dominated by the contributions of electrons below that level throughout, indicating that although the physics is more complicated than the 1D absorption model in general that model remains the most important piece to understanding the overall absorption dynamics.

can see clearly that the current is made up of a hot forward current and a cold neutralizing return current. Looking at 2.33.c we observe that the heat flux is carried exclusively by the high energy particles; two beams from the two different acceleration mechanisms are also apparent here.

2.3.2 Plasma Evolution and Laser Absorption in the Presence of Preplasma

Until now we have considered plasmas formed by short laser pulses without prepulses impinging on solid density targets. For such scenarios, the plasma profile is a step-function




FIGURE 2.33: *Distribution function moments plotted as a function of the logarithm of the particle energy, and x_1 . a) Density (0^{th} moment) b) Current (1^{st} moment) c) Heat flux (3^{rd} moment). Comparing a) and c), the separation into a background (a) and a beam (c) is clear to see; that the beam is actually two superimposed beams, one at the quiver energy and much hotter but less dense one, can also be discerned. The hot forward and cold neutralizing return currents can be observed in b).*

from vacuum, or a step function with an underdense shelf. Perfectly sharp plasma surfaces such as we have been discussing so far are a special case. In general a plasma will have a finite scale-length over which it goes from vacuum (or at least highly under-critical) to highly over critical density. This may either be because the target was assembled by compressing a lower-density target using an ablation process, the laser pulse is very long, or because the laser pulse also includes a low-intensity but long ‘pre-pulse’ before the main pulse which heats and expands the target surface.

Very high contrast ratio lasers (i.e. those with insignificant pre-pulse) do exist, and more can be expected to be developed in the future. Using such a laser to illuminate an initially solid-state target could plausibly create only a sub-wavelength scale preplasma, which would make the results of the previous sections directly applicable as is. We note that the previous results encompass the case where the main pulse itself rapidly generates an underdense region in front of the target self-consistently. However a more complete picture can be found by exploring what happens when the density scale length is initially longer than the laser wavelength. The goal is to identify how the physics of the idealized case relates to the more

general scenarios.

We find perhaps surprisingly that as before the main absorption mechanism which generates modestly energetic electrons remains the same as for a plasma which evolves from an initial vacuum interface - which is itself the same as the absorption mechanism from that case in early time, when the interface is still a sharp gradient to vacuum. For modest density scale lengths, this mechanism retains roughly the same efficiency as in the no pre-plasma case; although the generation of high-gamma particles from the low density shelf region is somewhat more dominant in the case of finite initial scale-lengths, resulting in greater overall absorption but also sometimes diminished absorption at the surface. These features are themselves all consequences of the fact that the response of the bulk plasma (as opposed to the accelerated electrons) is very similar either with or without a realistic density gradient - an underdense region transparent to the laser escapes freely back towards the direction of laser irradiation, while the critical surface is pushed inward by the laser pressure so that the laser reflection region is much smaller than the laser wavelength.

We model targets with finite density gradients as having a simple exponential dependence of density on longitudinal position with various scale lengths (we define the scale length as the distance over which the density changes by one decade.) The density ramps from underdense (and fewer than one particle per cell) up to 100 times critical density, after which it remains almost constant save the slight modulation to remove spurious resonant effects with the laser frequency. Otherwise the initial conditions of the plasma and the numerics are the same as those described in the previous section, being 3 microns wide and periodic in the transverse direction and having 80 microns of plasma in front of a particle damping region acting as a boundary condition.

Figure 2.34 shows the evolution of the density profiles under laser illumination for three different scale lengths: $2\mu m$, $3^{1/2}\mu m$, and $5\mu m$. As can be seen in the plots in the left-hand column, which shows the density in the entire initial vacuum region of $160\mu m$ with a logarithmic color axis, in all three cases an underdense plasma of a few percent of the critical density expands rapidly outward to a length much greater than the initial scale-length. The figures in the right hand column show the initial density evolution near the surface on a




FIGURE 2.34: Transversely averaged density for plasmas with different scale length pre-plasma ($2\mu\text{m}$, $3.5\mu\text{m}$, and $5\mu\text{m}$, top to bottom), shown logarithmically for the entire vacuum region (left column) and linearly as a relief map zoomed around the surface. The log plots show the expansion of an underdense plasma shelf into the previous vacuum region in all cases though quicker for longer scale lengths. The linear relief plots show the inward compression of the surface (“snowplowing”), which happens more slowly and to lower peak density in the presence of longer scale length pre-plasma.




FIGURE 2.35: Transversely averaged density around the critical surface (left column) and forward heat flux energy spectra (right column) as a function of time for three different pre-plasma scale lengths ($2\mu\text{m}$, $3.5\mu\text{m}$, and $5\mu\text{m}$, top to bottom). Bursts of near critical density plasma away from the surface can be seen to be associated with greater heat flux carried by high gamma particles.




FIGURE 2.36: Heat flux spectrum for 4 different scale length pre-plasmas (vacuum, $2\mu\text{m}$, $3^{1/2}\mu\text{m}$, and $5\mu\text{m}$), but now against the logarithm of the kinetic energy. As was also seen in figure 2.35, the heat flux carried by electrons accelerated at the plasma surface is not greatly affected by the presence of a modest pre-plasma; however the increase in the total heat flux carried by high gamma electrons coming from the underdense region can more clearly be seen in this figure.

linear z-axis (here both the color axis and redundantly the vertical axis of a 3D plot); again all three cases show similar behavior to each other, that is a steep sharpening to a sub-micron density scale length. Comparing to figure 2.30, we see a similar evolution in the case with an initially vacuum interface, to a long scale-length underdense region and a very short scale-length transition to highly over-critical densities.

These figures show that the transverse average density profile of the plasma for different initial scale-lengths at the surface evolve quickly under laser irradiation to very similar profiles both to each other and to the initially vacuum case; however when looking at the accelerated electrons it can be seen that despite the seemingly similar densities, these initial conditions do modify the absorption much later in time. Figure 2.35 shows the spectrum of forward heat flux, similarly to fig. 2.32 for a vacuum interface, along with a logarithmic plot of the density near the surface. There is somewhat more heat flux being carried by electrons above twice the quiver velocity than was seen in the case of no preplasma; that absorption especially happens in bursts associated with the presence of higher-density plasmas in front of the reflection region, as can be seen by comparing the spectrums with the associated density plots. However in these graphs otherwise the absorption spectrum seems roughly similar to the vacuum interface case; specifically the surface-accelerated electrons in the gamma range up to $2a_0$ are present as a strong beam in all cases, though the beam is strongest in the case of no pre-plasma.

Looking at the same heat flux spectrum data but against the logarithm of the kinetic energy, as shown in fig. 2.36 (along with the no pre-plasma case), we can see more clearly the enhanced high-gamma absorption in simulations with pre-plasmas. The beam of electrons accelerated at the surface can be seen clearly here in these figures as well, appearing as a band between kinetic energies of about 1.5 and 5 MeV. However the electron heat flux above this – coming from the stochastic acceleration in the underdense region – which appears as monotonically decreasing on a linear scale appears instead as a second peak here, centered roughly between 20 and 50 times the rest energy. This second beam is somewhat visible though weak in the initial vacuum case, excepting one period early in time where the beam is strong. In comparison it is enhanced in all the pre-plasma cases, though here too it is not a

constant feature, enhancing and depleting throughout all the simulations; the greatest high gamma heat flux appears in the longest pre-plasma, however the correlation to initial scale length seems somewhat weak as the $2\ \mu\text{m}$ simulation also shows a strong beam.

So, to summarize, the presence of a modest pre-plasma appears to noticeably increase high energy electrons coming from the underdense region but at the same time to have only a weak affect on the electrons accelerated at the surface. The tolerability, or desirability, of pre-plasma therefore depends on the goals of the experiment. If the desired outcome of the laser-plasma interaction is simply maximum coupling of the laser energy to the electrons, a thin pre-plasma may actually be desirable as it increases total heat flux. Or, if the goal is to maximize the heat flux in moderate gamma electrons (as is the case in Fast Ignition, where 1-3 MeV electrons are thought to couple most efficiently to the core), then some preplasma may not necessarily be helpful but it will at least be tolerable, as it doesn't much affect the surface acceleration mechanism. However, if very high energy electrons are in fact detrimental to the application, it would be optimum to keep any underdense regions to a minimum, both during the interaction and seemingly even before the main pulse.

2.3.2.1 Longer preplasmas

The cases considered so far cover cases up to only up pre-plasma scale lengths of $5\ \mu\text{m}$, which is rather modest. Although such scale lengths are plausible for a laboratory plasma, longer density scales are also likely to be encountered, especially in the case of an ablatively compressed plasma. The reason we have emphasized the results from the smaller scale-length simulations is that we have found that – with this simulation geometry at least – for longer scale-lengths the absorption can realize qualitatively different outcomes, depending chaotically on the initial evolution of the plasma. Whether these results can readily generalize to a 3D, finite spot-size target is not fully clear, but we include a discussion of them because of the interesting implications if they could.

This result is summarized in figure 2.37, which shows the heat flux spectrums for three slightly different simulations all with $10\ \mu\text{m}$ scale density gradients. The figure on the upper

left shows the spectrum for a simulation identical to those just described except for the ramp region. As can be seen the first phase of the simulation shows qualitatively the same behavior as before, with absorption both in the low gamma electrons coming from the surface acceleration and high gamma electrons coming from the stochastic underdense acceleration; quantitatively, that the high gamma beam is somewhat stronger and persists somewhat longer than for instance the $5\ \mu m$ case is also consistent, as it would take longer for the shelf to compress and become depleted.

However, after $t = 3000/\omega_0$, a completely different behavior appears: the heat flux from electrons in the surface acceleration energy band disappears almost completely, with the heat flux now being carried by the higher gamma particles exclusively. Looking at the density profile near the surface, displayed on the upper right of the figure, an expansion of the plasma into the underdense region can be seen to be coincident with this seeming decrease in surface absorption. It was therefore supposed that an enhancement in the density of the shelf region was leading to greater underdense absorption and subsequently to pump depletion before the laser could reach the surface region. (It's worth noting however that a plot of forward heat flux vs. position, not shown here, indicated contradictorily that absorption was happening at or near the surface).

The question then was what was leading to the growth of the underdense region. Of course it was possible that it was a physical effect, with the longer density scale-length leading to a different evolution of the multi-dimensional surface geometry, which allowed more electrons to escape through the reflective surface fields. But we postulated that it may also have been due to an artificial refluxing effect; specifically, we knew that the particle absorber at the rear of the target would re-emit particles with a temperature much greater than the background temperature, because when it measures the local temperature it makes no distinction between background and beam electrons and so most of the energy measured comes from the beam electrons. This could have been a problem for all the simulations using the absorber; but we thought what might have happen is that in the case of the longer density region, there would be a greater flux of high gamma electrons right at the start of the simulation, which would then ‘overload’ the absorber and lead to an initial burst of




FIGURE 2.37: The forward heat flux spectrum for three different simulations, all with $10 \mu m$ density scale lengths. The upper left figure shows the spectrum for a simulation using the standard particle damper to reduce refluxing; the upper right figure shows the density near the surface for the same simulation. The lower left figure shows the spectrum when using a damper which calculates the temperature in the damping region using a fourth-root weighting for the particle energy, to reduce the effect of the hot tail and so further reduce refluxing. The lower right figure shows the spectrum for a simulation using the same 4th root calculation, but also re-emitting the damped particles with a temperature further reduced by a factor of two from the temperature calculated.

refluxing - which would then reflux back into the laser and lead to an enhanced high gamma absorption, and so on.

To test this theory, we set the absorber to calculate the temperature using a sum of the fourth root of the particle kinetic energy instead of a linear sum, making the calculation less sensitive to high-energy tails (c.f. appendix B, especially B.3). The results of using an absorber with the modified sum calculation are shown in the bottom row of figure 2.37. The left-hand figure shows the result of the regular 4th-root algorithm. The right-hand figure shows the result using the same 4th-root sum in the calculation, but also reemitting the damped particles from a cooler distribution function with T reduced by a factor of 2; this risked resulting in an undesirable cooling at the back of the box, but as it would be even more aggressive at eliminating any refluxing were that to be the issue it would make diagnosing the problem more clear.

As can be seen, the results are not consistent with the refluxing hypothesis. The simulation with reemission at the calculated temperature does indeed show a return to the behavior of the shorter scale-length pre-plasmas, where the initial density gradient leads to high-gamma absorption at early times but surface absorption is present throughout. However the simulation with the extra cooling of the reemitted particles sees a return of the anomalous behavior with the unmodified damper, in fact even to a greater degree than before. We conclude from these results two things. One, there is no reason to suspect that our particle damper is not functioning properly. And two, based on that, and based on the fact that none-the-less changing the damper parameters leads to a drastically different response of the plasma, that response arises in a chaotic fashion; although we had intended to test the boundary condition, what we actually accomplished was modifying the seed noise of this chaotic event.

At this point we compared other diagnostics from the three simulations, to see what events correlated with the suppressed low-gamma absorption. We found that the suppressed absorption was associated with the generation of a strong B-field region just inside the plasma: the presence of this B-field leading to the elimination of low-gamma electrons, it's absence leading to 'regular' behavior of the plasma. This is illustrated in figure 2.38. The left-

hand column shows the time-averaged magnetic field and density for the same simulation as [2.37c](#); the right hand column shows the same data for the simulation represented in [2.37d](#); the top two rows show the density and magnetic field at the end of the simulation, the bottom row the transversely-average B-field for the whole duration of the simulation. Looking at the 2D field plots, although a static magnetic field of strength on the order of that of the vacuum laser field does appear in the simulation where surface accelerated electrons are not repressed, it only does so in a localized region; moreover, by comparing the field plot to the density plot at the same time, it can be seen that the B-field here is formed near the surface region. In the simulation where moderate gamma electrons are absent, in contrast, the static field is both larger, filling the entire box transversely; and it appears not near the surface but within the plasma itself, extending to the point of peak density around $x_1 = 1010c/\omega_0$. Comparing the density plots to each other, they also show differences, with the case where surface electrons are seen in the spectrum having a sharp surface - which would be the correct geometry for the surface acceleration mechanism - whereas the case where surface electrons are suppressed has a region of 10-20 times the (non-relativistic) critical density about two microns wide at the target surface, behind which is a strongly compressed region up to $250 n_0$. Not plotted here, the pressure is also quite different between the two cases, with the low-density region in front of the magnetic field corresponding to a high-pressure band with peak pressures of $80\text{-}100 n_0 m_e c^2$; no such band appears in the simulation without the strong magnetic field.

What we believe is happening, therefore, is the following. During the initial stages of the laser-plasma interaction, strong quasi-static magnetic fields build up around the critical surface as the plasma is compressed by the laser; we observe these initial fields in all cases with pre-plasmas, more strongly in cases with greater scale-lengths. Such fields are often observed in laser plasmas [[93](#), [94](#), [95](#)]. Due to the symmetry of the simulation, neither sign of the B field is preferred (only the B_3 component is allowed in the 2D geometry), and filaments^{[27](#)} of both types arise. As time goes on and the plasma compresses, these filaments

²⁷ ‘Filaments’ in the sense that all components of a two-dimensional simulation are filaments (which go to infinity); in the diagnostics they appear more as ‘blobs.’




FIGURE 2.38: The depletion of moderate gamma electrons is caused by the formation of a strong magnetic field just inside the plasma. The left-hand column shows the case where no field is generated, corresponding to the spectrum in the lower-left graph of figure 2.37; the right-hand column shows the generation of a strong B -field, corresponding to the lower-right graph of the same figure.

move and diffuse with the flow of the material; if filaments of opposite sign merge they annihilate, but if two filaments of the same sign meet they can combine into one stronger filament. If a strong enough region of magnetic field is able to form from these mergers, it can become self-sustaining, gaining the energy it needs to overcome magnetic diffusion from its interaction with beam electrons. This magnetic field then pushes back both on the bulk plasma and the beam particles (although the energy becomes so high in this region the distinction is hard to make); resulting for the former case in a high-pressure region between the field and the laser, a lower-density shelf pushed out by this pressure, and a strong build-up of density at the edge of the magnetic field, all of which are apparent in figure 2.38; and for the latter case in the inability of electrons which have been accelerated once by the laser at the surface from entering the material - the heat flux which does enter the plasma is either carried by particles accelerated in the underdense region, or by those which have interacted with the laser multiple times.

At the moment it is not clear how to control the growth of these magnetic fields; they simply appear or not in our simulations in the presence of pre-plasma based on the random orientations of the initial magnetic fields. Nor is clear from this simplified and highly symmetric geometry if these fields would appear or appear as readily in three dimensions, or in finite-size laser-plasma interactions. However they are interesting because, if they could be controlled, the generation of magnetic fields through the interaction of the laser with a density gradient may be a method for further controlling the characteristics of the accelerated electron beam. A more detailed examination of the interaction of high-intensity lasers with density gradients is warranted.

It bears emphasizing that all these simulations explore the absorption in the presence of preplasma in rather simple cases. The density gradient is modeled as perfectly exponential, and likely more importantly the laser is described as naïvely as possible, with no temporal (nor of course spatial) structuring whatsoever. The main point of these simulations is to show that the natural response of the plasma - rarefaction of the underdense region and steepening of the critical surface - is such that direct acceleration by the laser field at the plasma surface is generally applicable, and not just a special or idealized case. The results therefore should

be taken as more qualitative than quantitative (when a more general approach is warranted), and even the decreased coupling to low-energy electrons in the long scale-length case should be seen as instructive rather than predictive of the necessary outcome of an experiment. It is highly likely that a more tailored approach could further optimize the photon-electron coupling to a desired outcome - a point we touch on briefly in the next section.

2.3.3 Circular and Elliptical Polarization

As we have seen, the natural response of the plasma to picoseconds of laser illumination – whether starting from a perfectly sharp surface or with a modest density gradient – is to form two regions: a sharp gradient from underdense to highly overdense, and in front of this a long underdense region with a very shallow gradient. This results in two absorption mechanisms, the $2a_0$ mechanism at the sharp surface and stochastic acceleration in the underdense plasma. The stochastic acceleration gives rise to electrons with energies much higher than those accelerated at the surface, and so this at first might seem deleterious to something like Fast Ignition, for which only electrons with energies in a certain (lower) range are preferred. However it turns out it really doesn't matter because, somewhat surprisingly perhaps, the total energy input into the electrons with lower relativistic gamma factors isn't much affected by the laser absorption in the underdense shelf. Nonetheless it would be interesting and it might be useful if we could avoid the growth of the underdense plasma shelf, and it turns out we in fact can, by changing the laser beam from a linear to a circular polarization.

As we saw in section 2.2.5, for short laser pulses with circular polarization ($\lesssim 100$ laser periods) although there is surface heating due to the slipping stream instability in this case electrons are not accelerated by the $2a_0$ mechanism and total laser absorption remains low. This is a result of the structure of the laser field standing wave – the magnetic field anti-node everywhere at the plasma surface rotates in direction but stays constant in time, so no matter when (or where) an electron reaches the surface it sees a reflecting magnetic field and is never able to enter the laser electric field and gain energy. However it turns

out that, similar to the case of linearly polarized light, the surface does not stay smooth indefinitely; rather it becomes ‘bumpy’, with a characteristic scale length on order of the laser wavelength. Interestingly, although there is surface deformation in both cases, the effect on the plasma profile and hence the electron acceleration is qualitatively different between the two polarizations.

The plasma critical surface is shown in figure 2.39 after approximately 2 picoseconds of laser illumination for three different laser polarizations. In the top frame the laser has a normalized vector potential of $a_0 = 6$ and a linear polarization (and it is p-polarized, that is the electric field is in the plane of the simulation). In the bottom frame the laser is circularly polarized, with the same time averaged laser intensity as in the top frame. The middle frame shows the results for an elliptically polarized laser, where the ratio of the peak electric field in and out of the plane is 2:1 respectively, and again the time averaged intensity is consistent with the other two frames. In the top frame the surface can be seen to be ‘spiky’, with five or six sharp features pointed out almost directly into the underdense region. This is in contrast to what is seen in the bottom frame, where although the surface is also not flat, there is an absence of the sharp protrusions seen in the top frame; the shape is more ‘bumpy’ than ‘spiky’. The middle frame has a somewhat in between shape to the critical surface, with bumps but also a few shallow spikes which point at an angle and are not normal to the surface. Related to the surface shape we can see in the top frame an underdense region in front of the high density plasma, which as we saw in section 2.3.1.1 is created by plasma leaking out along these peaks. In the bottom frame no such underdense region is present, indicating that without the sharp structures the plasma is not able to pierce the standing wave magnetic field and allow electrons to leak out into the initial vacuum region. In the middle frame a much lower density shelf region is vaguely visible, indicating that the surface peaks do allow at least a few electrons to leak out.

The left-hand column of figure 2.40 shows the transversely averaged density as a function of time, for circular and elliptical polarization, and also for a circularly polarized laser and plasma with a $2\mu\text{m}$ initial density gradient. The top frame, for circularly polarized light and a sharp gradient, shows that throughout the entire simulation there is essentially never any




FIGURE 2.39: The plasma surface for three different laser polarizations, after ~ 2 ps of laser illumination. In (a) the laser has linear (p) polarization; in (b) the laser is elliptically polarized with a 2:1 ratio between the electric field in and out of the plane, respectively; in (c) the laser has circular polarization. In the top frame ‘spikes’ can be seen at the plasma surface, which are diminished in the middle frame and absent in the bottom frame. The underdense shelf, caused by these spikes, can be seen clearly in the top frame, vaguely in the middle frame, and it is absent in the bottom frame.

underdense plasma formed in front of the target. For the elliptically polarized laser beam, a thin underdense region can be seen to form and grow over time. However it is much thinner and lower density than was seen for linearly polarized light, having a 20 micron thick region of about 1% critical density and perhaps a 2 micron thick region at closer to 10% of critical density. (Recall for linearly polarized light that by the end of the simulation the plasma shelf reached the edge of the box, which is 160 microns from the surface.) When an initial density gradient is present an underdense region is observed throughout the simulation even when the laser is circularly polarized. The density of this region appears to decrease constantly as it expands outward though, indicating that it is not being fed by electrons escaping the bulk region.

The right-hand column of figure 2.40 shows forward heat flux as a function of particle γ factor and of time, as measured in a region just inside of the plasma bulk as described earlier, for the same three simulations as the left-hand column. The results are interesting. When using a circularly polarized laser, for the first roughly 500 fs, up to $t = 1800/\omega_0$ (the laser reaches the surface at $t = 1008/\omega_0$, the bottom of this figure), there is no laser absorption into energetic electrons, which is consistent with the results of section 2.2.5. However after this time, and even more so after $t = 2500/\omega_0$, the net forward heat flux goes up considerably – primarily into lower energy electrons, with energies a few times the rest mass, but with some energy going into electrons with up to twice the quiver proper velocity. Electrons with energies above this – the ‘superhots’ which are accelerated in the underdense shelf – are predictably completely absent. This indicates that although the bumpy surface is not able to fully pierce the standing wave magnetic field and allow electrons to escape outward, it does cause the surface field to distort sufficiently to allow electrons to enter the laser electric field and gain energy. In short, the electrons cannot get *through* the laser field but they can get *into* it. For Fast Ignition this is perhaps a very promising result, as electrons are being accelerated into, and only into, the proper energy range. Unfortunately it doesn’t actually appear to be very useful though, as there is no increase in total energy absorbed into lower γ factor electrons, as compared to linearly polarized laser light: total absorption is on the order of 18% for circularly polarized laser light, as compared to 28% of incident laser energy




FIGURE 2.40: Transversely averaged density (left column) and forward heat flux spectrum (right column) as a function of time for lasers with circular polarization (top row), elliptical polarization (2:1 polarization ratio, middle row), and circular polarization with a $2\mu\text{m}$ scale length pre-plasma (bottom row.) All cases show lower growth of a plasma shelf, and consequently greater coupling to lower-gamma electrons. In the case of pre-plasma, the shelf expands in space but doesn't grow in density, showing that the circularly polarized laser still contains most of the electrons in the bulk.

going into $2a_0$ electrons for linear polarization.

When an elliptical polarization is used (middle right frame of figure 2.40) absorption can be seen to increase in all particle energies, and consistent with the thin underdense region some superhot electrons can be seen, although they carry little total energy. The results are somewhat in between those of linearly and circularly polarized laser beams, which is perhaps intuitive as an elliptically polarized beam is a superposition of linearly polarized and circularly polarized beams. However importantly, just as was the case for pure circular polarization, there is no net increase in energy going into lower γ factor electrons, with roughly 25% of the incident laser energy being absorbed by $2a_0$ electrons. Finally, looking at the heat flux spectrum for a circularly polarized laser beam incident on a 2 micron density gradient (bottom right frame of figure 2.40), we see that other than a burst of superhots at the beginning of the simulation – due to the transient presence of a not insignificant underdense shelf – the laser absorption is consistent with the same polarization but a sharp density gradient. Indeed total absorption is actually decreased somewhat further, with roughly 13% of the incident laser energy being converted into energetic electrons. We find therefore that, although in contrast to the short pulse case ($t < 500fs$) laser energy absorption is possible with circularly polarized laser beams, in all cases total electron heat flux is diminished when using non-linearly polarized laser light.

2.4 Conclusion

We have looked at the absorption of high intensity lasers by near-solid density surfaces in one and two dimensions. We identified two mechanisms in our simulations, one taking place at the material surface, and one occurring in the underdense region in front of it. The former leads to an electron beam with energy bounded by twice the laser quiver velocity, the latter leads to a beam of unbounded energy; in the context of Fast Ignition only electrons within a certain moderate energy range are optimal for energy deposition, and so understanding and controlling the evolution of these two mechanisms is important.

We found that, although key portions of the surface absorption can be explained (and

modeled) in a single spatial dimension, those key portions are not actually complete by themselves, and so the total absorption phenomena can only be understood in a multi-dimensional framework. Specifically, we found that the standing wave structure caused by the reflection of the laser is critical to how electrons are accelerated by the laser and hence how the laser is absorbed, and the interaction with this standing wave structure depended on the *transverse* temperature, which in one dimension is an initial parameter which cannot evolve self-consistently.

However, the one dimensional simulations were able to make clear why the transverse temperature was important. Electrons accelerated at the surface and sent back into the plasma *all* had finite transverse momentum. In 1D, transverse canonical momentum is strictly conserved. Since the electrons start inside the plasma where there laser has not penetrated and they end up back in the plasma where there is no laser, then their final transverse momentum is identical to their initial momentum. Therefore only electrons which start with a finite (large) momentum can gain significant energy from the laser. Through particle tracking, we found that electrons with finite transverse momentum are deflected by the surface magnetic field (in the standing wave the magnetic field has an anti-node at the surface) to be traveling almost normal to the surface, so that they can penetrate into the vacuum to be accelerated by the electric field of the standing wave (which is an anti-node a quarter wavelength from the surface).

In addition, we found that the (relatively) steady-state mode of the underdense absorption turns out to be dependent on the transverse shape of the critical density surface, and so is a purely multi-dimensional effect as well. We further found that even modest density gradients in front of the over-critical plasma can affect the ratio of moderate- to high-energy electrons for times on the order of picoseconds (the longest possible duration for an ignition beam being on the order of $10ps$); although the compression of the surface to a highly over-critical density means the (preferred) surface acceleration appeared even with density gradients.

In both the initial vacuum and modest gradient cases, we found the use of a circularly polarized laser beam greatly decreased the growth of the underdense shelf and so limited the high-energy electrons, although this also reduced the total absorption. The use of an

elliptically polarized laser beam seems to give a compromise between total absorption and electron beam temperature, a promising result.

CHAPTER 3

Modeling of Electron Transport Experiment on OmegaEP

In the previous chapter we studied in detail how an intense laser is absorbed into hot ($p_1 \sim 2a_0m_e c$) and a super-hot ($p_1 \gg 2a_0m_e c$) groups of electrons. In this Chapter we describe simulations of an experiment which studied the generation and transport of relativistic electrons generated by a laser hitting a solid surface. The experiment was conducted on the OmegaEP laser system – codenamed eXport-10B (for ‘electron transport experiment 2010, shot day B’) – which examined the transport of relativistic electrons, generated by a high-intensity laser–matter interaction, through a relatively high-density and low-temperature plasma (often referred to as ‘warm dense matter.’) The experiment was carried out by a team at UCSD, including collaborators from General Atomics and the University of Rochester. The basic premise of the experiment was to create a volume of plasma using the long-pulse system on Omega ($\tau \sim 10^{-9}s$, normalized vector potential $a_0 \sim 10^{-4}$) incident on an ablator target layer above a plastic foam, then to fire a short-pulse laser ($\tau \sim 10^{-12}s$, $a_0 \sim 1$) along a perpendicular axis onto a gold foil, in order to generate a beam of high-energy electrons. This generated beam of electrons was then allowed to transport through the pre-formed plasma volume, and their emergence from their transport path would then subsequently be imaged via K- α x-rays generated in a copper foil on the far side of the target (along the short-pulse beam path). In order to have a second experiment for comparison, the same basic setup was repeated, but omitting the long-pulse laser beam, hence looking at electron transport through an initially room-temperature foam.

A third experimental setup was also used, with out the long-pulse laser beam, a thinner solid-density transport layer, and a decreased short-pulse laser energy. The results of these

experiments are included below, but due to the many changed variables – most especially a lower short pulse laser energy, which corresponds to a lower laser vector potential, and a decreased mean electron energy to below relativistic speeds – we don’t believe the results are directly comparable and so did not focus our simulation campaign on them.

The primary results of these experiments is that for unshocked foam transport layers, the $K\text{-}\alpha$ signal was over an order of magnitude greater than the signal from the pre-ionized plasma case. In addition the signal was significantly spatially wider in the plasma case, with the signal being almost uniform across the $480\text{ }\mu\text{m}$ witness copper, as compared to the cold foam with a localized $\sim 200\mu\text{m}$ wide peak. The inference from these results then is that the transport is significantly impaired in the plasma case compared to the cold case.

We have attempted to elucidate the mechanism(s) responsible for these results using Particle-In-Cell simulations run with OSIRIS. We have not simulated the long-pulse laser interactions with the ablator and formation of the plasma by the resulting shock, as these occur in a hydrodynamic regime not well suited to an explicit PIC code such as OSIRIS (radiation hydrodynamic simulations done by our collaborators will be discussed). Instead, we will focus only on the high-intensity, short-pulse laser-plasma interactions. As the reader may have already inferred, even this part of the experiment contains a rather wide range of parameter regimes for both the gold and the foam. The initial setup alone can span from solid-state high density metals to the moderate plasma state of ‘warm dense matter’. Furthermore, once the short-pulse laser is incident on the target, the plasma parameters can span even more orders of magnitude with the inclusion of a relativistically traveling electron beam. Due to the great breadth of physics potentially involved, it was quite a challenge to formulate questions which could readily be answered using the OSIRIS framework. None the less, through a non-trivial amount of time and effort, we have been able produce evidence for a plausible explanation for the experimental results.

The mechanism which the simulations indicate is the primary reason for the experimental result, is a partial resistive collimation of the forward-going electron beam in the initially unshocked foam, an effect which is absent in the warm dense matter. The resistive collimation itself is caused by the difference of the importance of Coulomb collisions between the two

cases. These differences are primarily (in fact, almost entirely) due to the density differences between the two cases. A secondary difference is an insulating magnetic field building up at the Au-CH interface in the case of low density plasma, caused by the convective transport of the magnetic field by the return current. This limits the heat flux from the gold into the foam.

3.1 Experimental Setup and Results

The geometry for the experimental setup is shown in figure 3.1. From the main diagram and the inset, it can be seen that when traveling along the path of the the long-pulse drive laser, the target consists of a plastic ('CH') ablator, a CH foam which acts as the transport layer, and a solid CH container. From the diagram, it can also be seen that when traveling along the path of the short pulse laser, the target consists of a gold foil to act as the electron source, the CH foam transport layer, and a copper foil to image the hot electrons as they exit the foam.

3.1.1 Target design

The Au foil was $10\mu m$ thick, and wider than both the laser spot and the CH foam region in the dimensions transverse to the short pulse laser beam path. The Cu layer was $25\mu m$ thick, and $480\mu m$ in both transverse directions. The CH foam had a density of $0.2g/cm^3$, was $250\mu m$ thick (i.e., $\rho r = 0.005g/cm^2$), and was $250\mu m$ wide in the horizontal direction and $400\mu m$ tall in the vertical direction. The foam was contained within a C-shaped solid plastic polyimide block, $1mm$ wide in the horizontal direction and $1.2mm$ tall in the vertical direction (and with the same thickness as the foam), and with a density of $1.4g/cm^3$. The front of the C-shaped block, along with the surface of the foam, is covered with the ablator, which is a $40\mu m$ thick polypropylene foil ($0.91g/cm^3$).

In order to shock the foam, the ablator is irradiated with 351 nm UV light ('3 omega'), with an energy of 1.2 kJ in a pulse length of 3.5 ns and a spot with a diameter of $150\mu m$. This corresponds to a peak laser intensity of approximately $I \sim 5 \times 10^{15} W/cm^2$. The an angle




FIGURE 3.1: *Target setup and geometry for the eXport10B experiment.*

of incidence was 24° . In all transport experiments the gold foil was irradiated with 1053nm IR laser light containing 760J in 8.1ps at an angle incidence of 25° . The peak intensity was estimated to be between $4 \times 10^{18} \text{W/cm}^2$ and 10^{19}W/cm^2 , giving a spot-size of between $27 \mu\text{m}$ and $17 \mu\text{m}$.

As mentioned earlier, a third experiment on a second target was also fired, which consisted of a similar setup but with the foam region replaced by solid plastic. Also there was no ablator as the target was only used in the initially cold state when short-pulse radiation began, and the thickness of the target was reduced to $50 \mu\text{m}$, in order to retain the same areal density (ρr) as the cold CH foam. However this target was only used once, and that shot only had a total laser power of 1/3 of that used in other laser shots (266J in 8.5ps), making comparative analysis of the results difficult.

3.1.2 Plasma Characterization

3.1.2.1 Experimental Characterization

To characterize the CH foam plasma created by the shock launched from the ablator, a series of shots was carried out without the short pulse laser [97, 96]. In these shots, as in




FIGURE 3.2: The x-ray absorption spectrum for an experiment using the long pulse beamline only and an aluminum-doped foam (solid black line), along with simulated absorption spectra from the code FLYCHK (colored broken lines.) The x-ray absorption of the Al-doped plasma was found to be most closely consistent with a plasma having a temperature of 40eV and a density of 30mg/cm^3 . Figure from [96]

the transport experiments, the plasma was created by irradiating the ablator with a single Omega EP 351nm beam. Unlike in the transport experiments, the CH foam was doped with Al_2O_3 such that the density of aluminum atoms was $4 \times 10^{20}/\text{cm}^3$. Located 2.5mm away from the target, an $80\mu\text{m}$ Samarium (Sm) ‘dot’ was irradiated with three more UV beams, which acted as an x-ray light point source. On the far side of the shocked plasma, an x-ray streak camera was used to image the x-ray spectrum as a function of time. The x-ray spectrum from the Sm dot was also directly measured with a time-integrated Bragg crystal spectrometer. Comparing the original x-ray spectrum with the spectrum transported through the plasma, the absorption by the Al species was measured. This spectrum was compared with simulated spectra generated using the atomic physics code FLYCHK [98]. The spectrum at time $t = 7\text{ns}$ after the drive laser is turned off is shown in figure 3.2. As can be seen, a best fit was found to be a plasma having a temperature of 40eV and a mass density of $0.030\text{g}/\text{cm}^3$ ¹. The temperature and density were assumed to be uniform.

3.1.2.2 Radiation-Hydrodynamic Simulation

Radiation-hydrodynamic simulations using the code DRACO [99] were performed in 2D. Results show the plasma in the center of the foam region having a density n_e of $\sim 25 \times 10^{21}/\text{cm}^3$ and a temperature T_e in the range of $\sim 30 - 70\text{eV}$ at $t = +8\text{ns}$; these results are similar to the experimental results, though both are slightly higher. More importantly, the simulation shows a relatively uniform plasma in both density and temperature in the region behind the short-pulse laser focus.

3.1.3 Experimental Results of Electron Transport

The results from the transport experiment are summarized in figure 3.4, which shows data from both the Spherical Crystal Imager (SCI) [100] and the Zn von Hamos spectrometer (ZVH). The main result is a clear and repeatable decrease in measured K- α signal by greater

¹ Atomic absorption is sensitive to the concentration of atoms and hence indirectly to mass density. No diagnostics were performed which would be sensitive to plasma density, and so the charge state of the carbon ions is not directly known. Assuming the charge state of carbon is +4 (and fully ionized hydrogen) the electron density would be $6.9 \times 10^{21}\text{cm}^{-3}$.




FIGURE 3.3: *Electron temperature (top row) and density (bottom row) at three sequential times, as predicted by a simulation using the 2D radiation hydrodynamics code HYDRA. White circle indicates the diagnostics window from the experiment, black lines the boarder between the material from the initially foam and solid plastic regions. At the latest time, the density is uniform in the diagnostics window, though the temperature ranges from 30-70 eV. Figure from [97]*

than an order of magnitude in the driven foam case as compared to the cold foam shots.

Figure 3.4 b) and c) show the raw images captured by the SCI Imaging plates in the cold and driven cases, respectively; the gain has been set individually in each photograph, as otherwise nothing would be visible in the driven image. A central hot-spot in the cold case and basically uniform emission across the foil in the case of pre-formed plasma can clearly be seen. Figure 3.4 d) shows a central line-out of the same data, now with consistent normalization, with linear a y-axis and as a log-linear plot as an insert. The greater than 10-fold difference in signal intensity is apparent. Figure 3.4 e) plots peak SCI signal and total ZVH signal for two cold foam shots and three driven foam shots, showing the repeatability of the decrease in signal in the driven foam case. Figure 3.4 d) and e) also have results for the case when only the driver laser is fired, showing that at least 75% of the signal is from the influence of the the short-pulse laser. Also shown are results for a cold, solid CH transport medium, but as there is only a single shot and it is at a much lower intensity it's not completely clear how to interpret the data (all data have been scaled by the ratio of the short-pulse laser energies).

3.2 Simulations

3.2.1 Challenges and Limitations

As stated in the introduction, though superficially simple, this experiment is in fact complicated because it contains quite a wide range of parameter scales, inherently making modeling it a challenge for any simulation code, OSIRIS notwithstanding. We believe the choice of an explicit PIC code like OSIRIS was the correct one for this problem – as opposed to a more approximate code but one which could be used for longer time frames or higher densities for instance, such as an implicit PIC code or a radiation-hydrodynamic code – because the central facet of the interaction is the generation and transport of a relativistic electron beam, and this is inherently a highly kinetic and non-linear process. That it was the best choice however in no way means that it was a perfect one, and indeed the problem proved




FIGURE 3.4: Summary of the experimental results. 2D images of the $k\text{-}\alpha$ signal (a-c) and line-outs from the same (d) for the cases of solid foam, cold foam, and driven foam (note the color map is not to the same scale in the three 2D plots). The integrated total luminosity for the three cases, and driver only, is presented in (e), showing a greater than 10-fold drop in $k\text{-}\alpha$ in driven versus shocked foam.

impossible to answer using only a pure PIC framework; the challenge then was to find the simplest extension to PIC which was able to reproduce the experimental results.

One important extension ended up being simply the inclusion of binary Coulomb collisions – however more specifically the answer was a model which could simultaneously calculate the high collisionality of the cold background matter efficiently, and calculate the lower collisionality of the relativistically traveling beam correctly. Such a model was developed by Pérez in a paper published in 2012[60], and we implemented this into OSIRIS. OSIRIS already contain a collision framework including models based on work by Nanbu [56] and Sentoku [58], so this framework was extended to incorporate the Pérez algorithm.

Although we were able to include a consistent model for binary Coulomb collisions in our code, which as above stated we believe is key, there are of course other approximations or omissions which we still had to make to perform the simulations. We review some of these approximations now, so the reader can better understand the limitations of the results, which could then explain any deviations from the experimental data. Based on our understanding of these approximations, we incorporated some ad-hoc modifications to get better agreement. One limitation is that modeling the entire volume of the three-dimensional target was determined to be infeasible, and so a reduced-size and 2D system was simulated instead. We believe most of the relevant physics would be present in the 2D regime, but this did make direct quantitative comparison with the experiment difficult.

Another limitation is that the disparate scales between the main transport layer of low-density plastic and the electron ‘source’ layer of solid gold (and to a lesser extent the copper witness layer) cannot be modeled efficiently simultaneously in pure PIC. For instance assuming a charge state of $Z=40$ for the gold ions (which is roughly consistent with the energy density we see in our simulations in the entire volume behind the laser focus) is over 2300 times the laser critical density. In addition to increasing the total number of particles proportionate to the increase in density, this would also reduce the allowable time-step (as the plasma frequency must be resolved everywhere for stability), and the total increase in computation cost could be over a factor of 100, which was unfeasible. Much of the laser absorption and electron transport physics for highly relativistic electrons can be modeled using a lower

density plasma. We therefore chose to model the gold with the correct mass-density – so that the hydrodynamic motions would be reasonable – but with a plasma density (either 98 or $100 n_c$) corresponding to an average charge state a little bit below $Z=2$. This was deemed a reasonable approximation based on our previous work which indicates that so long as the target density is much greater than the laser critical density, the spectrum of accelerated electrons is more sensitive to the laser parameters and target geometry than on the exact target density.

However, the use of a lower electron density lowers the heat capacity of the gold layer. As a result, the background of the gold heats up to much higher temperatures than in reality. This heat is then transmitted to the plastic (Carbon). Therefore, both the simulated gold and plastic layers have artificially high temperatures. Most importantly, this causes the collisionality of the plastic to be significantly lower in the simulation than in the experiment. We believe that this is the main obstacle for getting quantitative agreement for the $K - \alpha$ between simulations and experiment. As we describe later we artificially increase the collisionality in the plastic (and gold) to compensate for the unphysically high plastic temperatures and this leads to dramatically different results.

Lastly, all the experiments were, somewhat by design, on the edge of the validity of the plasma assumption. In the case of the pre-shocked plastic foam, the plasma parameter – that is the number of particles in a Debye sphere – at the time of initial irradiation by the high-intensity laser was on order of 3; and in the case of cold CH it was of course 0 because the matter would be unionized at that point. This matters primarily for two reasons. First, even if the assumption that the system being modeled was a ‘plasma’ in the sense of being an ionized gas is valid – which would indeed be the case for warm dense matter – for PIC to be accurate necessitates that the system is a plasma in the sense of having a large plasma parameter, so that collective effects dominate over few-body effects. In this regime the results will be insensitive to whether there are few (PIC) or many (reality) actual particles interacting. PIC can be insensitive to the number of simulated particles because in PIC the collisions between particles are greatly reduced for impact parameters smaller than the particle (or cell) size. Also collisional physics becomes both more dominant and less well approximated by




FIGURE 3.5: *The initial electron density profile for the standard eXport10B simulation, for the low density CH case ($n = 7n_c$). In the medium density case the plastic region has a density of $28n_c$, in the high density case $63n_c$, with the geometry otherwise unchanged.*

the traditional model of small angle two body Coulomb collisions as the plasma parameter decreases. Secondly, and perhaps somewhat obviously, any physics which is strictly outside of PIC and traditional plasma physics and which becomes important at low temperatures – things like covalently bonded atoms and impact ionization – are fully absent from PIC and could at best be added in a very ad hoc manner (which in general we have chosen not to attempt).

3.2.2 Simulation Setup

The basic geometry for simulations of the eXport10B experiment is shown in figure 3.5. As can be seen, the target – which is isolated from the walls by a 10 micron gap on all sides (all units are normalized to the laser parameters) – consists of a 10 micron thick high-density region in front of a 100 micron thick lower density plasma; transversely the target is 80 μm

wide². A laser will be launched from the left-hand ($x_1 = 0$) wall and focused on the center of the surface of the high-density region, with a spot-size (gaussian full-width half-max) of $17 \mu m$ and a peak normalized vector potential of $a_0 = 3$.

The high-density region - which is meant to represent the gold foil from the experiment - has the same spatial setup in all simulations. The charge density for both species is 98 times the laser critical density, which is approximately the same plasma density as twice ionized solid gold; the charge-to-mass ratio of the ions - $5.6 \times 10^{-6} e/m_e$ - is also consistent with Au^{+2} , so that the mass-density of this region is comparable with the mass-density from the experiment, even if the charge-density is grossly incorrect. As noted earlier, this leads to the gold layer becoming artificially hot because of its low heat capacity (low electron density). When pertinent, for the purposes of Coulomb collisions the ions are also given a charge of $Z = 2$ for lack of a clearly superior choice, but as this is not expected to be physically correct we caution against drawing many conclusions from any collisional effects seen in the gold region. The initial temperature of the gold region depends on that of the plastic region, which is explained below. In cases where the plastic region is hot, the temperature of the gold region is set such that there is a pressure balance between the two, as otherwise the interface can be seen evolving in the short time before the laser has even reached the target. When the plastic region is at room temperature, the gold region is as well, as the pressures are trivial in this case. Electrons and ions are given the same temperature.

The lower density region is meant to represent the plastic transport layer of the experiment. Three different plasma densities were used: a low density of seven times critical density (for a $1 \mu m$ wavelength laser), which roughly the density of fully ionized 25 mg/cc plastic; an intermediate density of $28n_c$, which is equivalent to approximately 90mg/cc; and a high density of $63n_c$, which is equivalent to approximately 200mg/cc. The low and high density simulations are meant to correspond to the pre-shocked plasma and cold foam experimental cases, respectfully. Therefore in most cases, the low density case is initialized with the experimentally measured temperature of 40 eV, and the high density case with a temperature of

² The x_2 offset from 0 is not meaningful here; it's an artifact of being derived from a simulation with a wider geometry.

0.026 eV (i.e. room temperature). The intermediate density case is meant to imagine a lower density foam, and so in those simulations the CH region is also initialized with a temperature of 0.026 eV. For simplicity rather than a multi-ion species plasma, the plastic region is instead modeled as pure carbon; this leads to a charge-to-mass ratio of $2.74 \times 10^{-4} e/m_e$ and an ion charge state of $Z = 6$.

The numerical parameters are as follows. The PIC cells are square with a size of $0.15c/\omega_0$ on each edge, which gives a total box size of 5460×4096 cells. The time-step for the PIC loop is $0.05/\omega_0$ which is less than the Courant condition. This allows us to properly resolve the plasma frequency even in the high-density region, which we found to be necessary to avoid heating due to Coulomb collisions. Initially all of the electrons have the same charge, and hence the number of particles-per-cell (PPC) must be variable since the density is not constant. The normalization is such that one particle in a cell gives a density of $1.75 n_c$, which works out to 4 PPC in the low density plastic, 16 PPC in the intermediate density plastic, 36 PPC in the high density plastic, and 56 PPC in the gold. Electrons and ions are both initialized with the same number of particles (even though $Z=2$ for the gold region and $Z=6$ for the carbon region). To reduce the effect of macro-particle stopping (see Chapter 4), which otherwise would be the dominant stopping and heating mechanism for relativistic electrons with these parameters, splitting of the high energies electrons is used. The particles are split at kinetic energies between $0.4 m_e c^2$ and $1.0 m_e c^2$ and are split up to 4 times, for a reduction in macro-particle size by a factor of 16. We use cubic (3^{rd} order) particle shapes in order to effectively eliminate grid heating.

In addition to varying the density, we have run the simulations with different collisionality physics for the plasma. We have three collisionality cases, which we call the no-collisions case, the correct-collisions case, and the enhanced-collisions case. ‘No-collisions’ means that we run OSIRIS without any Monte Carlo collision model, so we are in a pure PIC situation (there are, however, still finite size particle collisions). ‘Correct-collisions’ means that we use a Monte Carlo binary collision algorithm where all of the collision parameters, including the Coulomb logarithm, are calculated automatically from the particle and fluid parameters (moments over the particle data). ‘Enhanced-collisions’ uses the same collision algorithm,

with the exception that instead of calculating the Coulomb logarithm automatically we fix it for the whole simulation at some high value, usually 50. Note that although fixing the Coulomb logarithm this way is of course non-physical, we argue it may lead to results which more closely agree with the simulations. As we noted earlier because we model a gold layer with a lower density the heat capacity is lower. This leads to an artificially high temperature for the gold and hence the plastic, which decreases the collisionality. We are attempting to compensate for this. Also note that, although 50 is a very high value for the Coulomb logarithm, which would be expected in the most extreme cases if ever, increasing the Coulomb logarithm is roughly equivalent to increasing the atomic number Z of the ion species (under the assumption that ion-electron collisions dominate the physics and the Coulomb logarithm does not vary with position) which is physically reasonable. However as our goal was not study the physics of different materials per se, but rather to look specifically at the effect of higher collisionality, we choose to focus on the Coulomb logarithm instead.

The Monte Carlo collision model which we used is that of Pérez et al.[60] including the low-temperature limit on the collisionality explained therein. By including the low-temperature cutoff, the collision frequency is prevented from ever exceeding the plasma frequency, even in very cold matter, and so we use a collisional time-step of twice the PIC time-step, $dt_{collisions} = 0.1/\omega_0$. The collision cell is the same as the PIC cell (although it need not be).

Longitudinally, the boundary conditions are absorbing for the fields (perfectly matched layer) at the front and back of the box, and specularly reflecting for particles. Transversely the boundary conditions are usually periodic, though sometimes the same conditions are used as the longitudinal boundaries (absorbing for fields, reflecting for particles) and the results are materially the same. Other than for the reflection of the laser, in general we would not expect much boundary effect in this simulation, due to the 10 micron vacuum gap on all sides and the relatively high mass density of the plasma, so the insensitivity to the transverse conditions is not surprising. In fact, we prefer the use of isolated targets for these very reasons.

3.2.3 Simulation Results

3.2.3.1 Base cases: Solid State versus Warm Dense Matter

First, we discuss the two cases which one might naïvely believe are closest to the ‘correct’ physics of the experiment, those being the high- and low-density plastic transport regions with Coulomb collisions on and with the Coulomb logarithm calculated automatically based on the local plasma temperature and density. Figure 3.6 shows the forward heat-flux for the two cases, at the end of the simulations, which is after ~ 4 ps of laser illumination ($\Delta t = 8000\omega_0^{-1}$, 4471 fs more exactly). Although the experimental K- α diagnostic is more directly related to heat than to heat-flux, we focus on the latter in simulation diagnostics because the low heat-capacity of the low density gold and 2D geometry is particularly detrimental to any sort of simulated K- α post-processing. The plasma becomes almost uniformly hot everywhere, which leads to near-uniform X-ray emission from the entire carbon region. The most striking feature of these figures is the filamentation of the heat-flux, and the difference in the filamentation between the two cases.

In the simulation with high density carbon (fig. 3.6.b), we see multiple filaments radiating out from the laser focus. The filaments travel in almost straight lines away from the laser absorption region, although there is some deflection and merging apparent. Most of the heat-flux is carried in filaments inside of a 45 degree cone from the focus point, though weaker filaments can be seen filling out a roughly 90 degree cone. Between the filaments, especially between the ones going almost straight, are regions carrying no forward heat flux at all, in fact they appear to be carrying a weaker return heat flux (light green color).

The heat flux for the low density carbon target appears quite different. Although filamentation is also apparent, instead of many filaments now there are just two main filaments traveling down the middle of the target. These filaments don’t go straight but bend and kink as they go through the carbon region. They also do not continue all the way through the target, fading away for $x_1 \gtrsim 400c/\omega_0$. Although these filaments carry the most intense flow of energy, the whole target shows a mostly featureless forward flow of heat.

Figure 3.7 plots the forward heat flux as a function of the longitudinal position and the




FIGURE 3.6: Forward heat flux (q_1) for plastic region densities of $7n_c$ (a) and $63n_c$ (b). The low density case shows a couple heat flux filaments, but they are crooked and don't seem to reach the end of the box, and they are surrounded by unstructured heat flow. The heat flux in the high density plastic has strong filaments which cross the box and in some cases converge with each other.




FIGURE 3.7: *The forward heat flux in the gold region as a function of forward position and electron gamma factor. This data is from the case of $63n_c$ density plastic, but as the gold layer is the same in all cases the diagnostics for all cases appear similar. Jets of electrons consistent with $2a_0$ acceleration are apparent.*

relativistic gamma factor of the particles in the front of the gold layer (averaged transversely across the entire target). The data shown is from the simulation with a high density plastic layer; however all cases appear similar (both the two we have mentioned so far and the different simulations we discuss later), due to the gold layer being similar in all cases. Jets of electrons can be seen going into the plasma spaced at half the laser wavelength (and hence entering at twice the laser frequency) and with peak γ of roughly $2a_0 - 1$, consistent with the $2a_0$ acceleration mechanism discussed in Chapter 2. Although over 4 ps of time has passed since the laser turned on, there are only a few ‘superhot’ electrons (those more energetic than the $2a_0$ electrons), a consequence of the high mass density of the gold ions allowing for only a shallow underdense plasma shelf to form in front of the target (the shelf is 3-5 μm thick and roughly 5% of the critical density).

Figure 3.8 shows the magnetic field for the same simulations and time ($t = 8000\omega_0^{-1}$) as figure 3.6. In both the high and low density cases, regions of relatively strong magnetic can be seen wrapping around the filaments of heat-flux. That the filaments stop short of the end of the box in the low-density carbon is perhaps more clear in this diagnostic; some additional,




FIGURE 3.8: Magnetic field (B_3) for $7n_c$ (a) and $63n_c$ (b) plastic layers. In both densities, the field can be seen to be wrapping around the filaments which were shown in figure 3.6 In addition in the case lower density CH there is a strong, turbulent magnetic field starting at the gold-plastic interface layer and extending into the plastic layer.

shorter filaments are also apparent in the magnetic field, though the field surrounding the main filaments seen earlier are notably stronger. One can also observe a region of very strong magnetic field in the low density carbon just outside of the gold, being roughly 10 microns thick and 30 microns wide; no such region is present in the high density carbon.

The right-hand side of figure 3.9 shows the transverse rather than forward heat flux for the two cases; again important differences are apparent. For the high density carbon case, the transverse heat-flux in the carbon region follows the same filaments as the forward heat flux, that is to say it is simply the transverse component of motion of the same particles which carry the forward heat flux; in the gold region the same holds, the energy from the laser absorption is carried in straight lines through the gold. In the low density carbon case, less transverse heat flux in the carbon region is visible than for the high density case, which is consistent with the fact that the although smaller filaments appear to be directed mostly forward in this case. However, this also means that the energy which heats the carbon doesn't travel either directly from the laser focus or latterly from the filamented region but rather forward from the gold region. Looking at the heat flux in the gold, we also see a strong difference (even though the gold is the same in both cases), with much more heat flowing transversely from the interaction region in the presence of the low density carbon.

The left-hand side of figure 3.9 shows the energy density or scalar pressure (temperature times density) again in both cases. In the high density carbon although the whole region is quite hot, some structure can be seen in the heat distribution on the right side of the box following the same filaments as all other diagnostics. Closer to the gold/carbon interface any structure is much harder to discern in this diagnostic. In the gold foil there is a very high pressure area directly behind the laser focus, but away from the focus the gold and carbon approach pressure balance. The situation is quite different with the $7n_c$ carbon plasma — there is very little structure visible anywhere except in the carbon region right behind the laser focus, where the two filaments can be seen coming out of a blob of high-pressure; otherwise almost the whole carbon region is of uniform pressure. The gold is now at a uniformly high pressure and at no location (even far from the laser focus) in pressure equilibrium with the carbon near it. We discuss the reasons for this later.




FIGURE 3.9: Energy density (pressure) in the left column and transverse heat flux (q_2) in the right column for $7n_c$ (top row) and $63n_c$ (bottom row) plastic layers. For the simulation with lower density plastic, the transverse heat flow is mostly confined to the gold region, and much of the heat ends up trapped in this layer. In the presence of higher density plastic, there is no significant heat flux laterally in the gold and much less energy is contained there.

It is also important to realize that, although the pressure is lower for the low density plastic compared to the high density plastic, because of the lower density the temperature is actually higher in the former case. Dividing the energy density by the plasma density one arrives at a temperature of 60 keV in the low density plastic, 8 keV for the higher density; however this also includes the energy of the relativistic electrons, and so exaggerates the temperature in each case. Measuring the distribution function we estimate a temperature of the background electrons of 12 keV in the low density plastic, 3 keV for the higher density. Therefore the low density plastic has a lower collisionality, and hence higher conductivity as conductivity is independent of density.

3.2.3.2 Comparison of Density and Temperature Effects on Filamentation

As explained earlier, to be consistent with the experimental parameters, both initial temperature and density were varied in the figures of the previous section; the $7 n_c$ plasma was initialized with a temperature of 40eV, and the $63 n_c$ plasma was initialized with the temperature of 0.026eV, i.e. 300°K or room temperature. However, as can be seen in figure 3.10 – which shows the forward heat flux at the end of each simulation for all four combinations of initial temperatures and densities (repeating two graphs from earlier) – the initial temperature doesn't seem to have much of an effect on the heat transport in the simulations; other diagnostics examined show similar consistency between the initially warm and cold cases. The low density case, in the left-hand column, does show a small but interesting variation when the temperature is reduced, in that there is now one main filament shooting off to the side of the target. However, most of the heat-flux is still carried in an unstructured way, and this feature may just be a statistical variation. The high density case shows even less variation, with perhaps slightly straighter filaments in the warmer case.

The insensitivity to the initial temperature is likely caused by the rapid heating of the background gold plasma from the high-intensity laser itself. The gold then heats the plastic. In part, this is certainly exaggerated by the incorrectly low heat capacity of the simulated gold layer which was mentioned earlier. However we argue that the unimportance of the




FIGURE 3.10: *Forward heat-flux compared for carbon transport regions at densities of either $7n_c$ or $63n_c$ (top and bottom rows, respectively,) for initial temperatures of 0.026 eV ($300^\circ K$) or 40 eV (left and right columns, respectively.) Density has a very strong effect, initial temperature near none at all.*

initial temperature is still plausible for the experiment. The reason is that 40 eV is still a very low energy density compared to the amount of energy in the short pulse laser. For instance, in the coolest part of the gold in the simulation with the high density carbon, the energy per gold atom is approximately 20 keV. If we assume the gold is at equilibrium at a single ionization state, and we use a published list of ionization energies for gold (for instance Beiersdorfer et al. 2012 [101]) we arrive at an estimated temperature in the gold of 660 eV (and a charge state of 23). It is not likely that the ionization state of the gold will have an ionization energy greater than the local temperature, so if anything this estimate is low. How and how quickly this heat will move into the plastic region is unclear from this simple estimate, but that the plastic will reach 10% of this temperature seems reasonable.

In any case however it is certainly true that in our simulations the initial temperature does not play a role. As such, although in all simulations (besides the two we just spoke of) we will initialize a warm plasma in the low density carbon case and a cool plasma in the high density carbon cases, we will not focus on the initial temperature any further.

3.2.3.3 Varying Collisionality

Next, we look at the result of varying the effect of collisions in the plasma, which as explained earlier will involve three cases for each density: a ‘no-collisions’ case, a ‘correct-collisions’ case, and an ‘enhanced-collisions’ case.

Let us start by looking at the effect of collisionality in the high density $63n_c$ carbon plasma. The forward heat flux and magnetic field strength for that density, after $4ps$, is shown in figure 3.11. The three cases show clearly distinct behaviors. In the collisionless simulation, other than some transverse localization in and near the gold, the heat-flux basically fills the box uniformly. Similarly the magnetic field, although present turbulently at the Au-CH interface behind the laser focus, is (uniformly) nearly absent throughout the carbon region.

This is in contrast to the results obtained by enabling the collisions. The heat-flux is transformed from uniform to highly filamented and directional. Although somewhat weaker filaments can be seen filling out a cone of perhaps 90° , most of the heat flux is within a

45° cone pointed straight forward from the laser focus. Furthermore the filaments, once they emerge by $x_1 = 200c/\omega_0$, travel in almost straight lines, although some merging can be seen (in a spatial, not temporal sense.) Looking at the magnetic field reveals the same pattern as seen in the heat flux, with strong fields wrapped around each of the heat-flux filaments. Although not shown here, when either of these diagnostics are looked at temporally, the basic filamentary structure can be seen arising early in time and then remaining roughly fixed in shape while growing in strength.

“Artificially” fixing the Coulomb logarithm leads to another dramatic change in behavior. In the heat-flux, although some similar filamentation can be seen down the middle of the box, the more obvious feature is that all of the heat-flux has been confined to a rather narrow forward-going beam. Indeed not only is the initial angle of divergence of the heat flux modest, but further into the plastic the heat flux filaments actually appear to turn and *converge* towards the forward direction. The magnetic field plot shows a filamentary structure along axis which is much clearer than can be seen in the heat flux. Furthermore, one can see one a strong magnetic field wrapping around the whole heat flux beam. The edge of the heat flux follows these magnetic field contours.

The same q_1 and B_3 diagnostics for medium density $28n_c$ carbon plasma can be seen in figure 3.12. Broadly these simulations follow a similar trend to the high density carbon ones, though there are differences. In the case without collisions, the heat-flux again shows little structure and is roughly uniform in the transverse direction. In this case the magnetic field is also nearly absent in most of the carbon region; although the turbulent region around the material interface is somewhat larger and stronger and some short filaments enter into the plastic. Turning collisions on again creates a filamented pattern in the heat flux, which again is seen in the magnetic field. However, in both cases the pattern is somewhat weaker with this lower (medium) density, and the size of the filaments is smaller and their total number is greater. Using a fixed Coulomb logarithm of 50 with medium density carbon leads to a somewhat interesting behavior, which upon further inspection can actually be somewhat discerned in the high density case. That is that a bifurcation of the heat flux into two forward beams can be seen. This aside, the behavior is otherwise fairly similar to what




FIGURE 3.11: *Forward heat flux and magnetic field in high density plastic with either collisions turned off (left-hand column;) with collisions on with the Coulomb logarithm calculated automatically (middle column;) or with collisions on and with the Coulomb logarithm set to 50 (right-hand column.) Without collisions the magnetic field in the plastic remains near zero and the heat flux has essentially no structure to its flow. With the automatic Coulomb calculation a strong and filamented field appears, and heat flux now follows that field. With $\ln(\Lambda_C) = 50$ the field filaments become surrounded by an envelope field, and the heat flux is constrained entirely within this envelope*




FIGURE 3.12: *Forward heat flux and magnetic field in medium density plastic with varied collisions, arranged as in fig. 3.11. Some weak magnetic field structures can be seen in the plastic region near the gold surface, but the heat flux remains mostly structureless. Using the automatic Coulomb logarithm creates filaments in both B_3 and q_1 , as in the higher density case, though both effects are somewhat weaker here. Setting $\ln(\Lambda_C) = 50$ again creates an envelope field, though interestingly one with a ‘split’ down the middle.*

was seen before. There are some filaments in the heat flux, but globally all the energy is beaming forward (nothing hits the side boundaries). Filamentation is more apparent in the the magnetic field but there is also a field wrapped around the entire beam(s). The initial divergence of the heat flux is greater than it was in the high density case, and although the beams do appear to start curving back towards the end of the box, the total width is notably broader.

The results for low density $7n_c$ carbon plastic, shown in figure 3.13, are significantly different. The $7n_c$ case is very distinct from both the $28n_c$ and $63n_c$, which were qualitatively similar to each other. For $7n_c$ the behavior is insensitive to the collisionality. For the collisionless case in the left-hand column, interestingly some filamentation is now apparent in the heat flux diagnostic. There are two ‘main’ filaments going roughly down the middle of the

box, and other weaker filaments going at odd angles. Although not shown, an examination of the temporal evolution of these filaments shows there is significantly less filamentary structure early in time than in the high density cases, and what structure does appear is more inclined to evolve and move about (rather than reach a quasi-steady state as was the case at higher densities). Consistent with the filaments in q_1 , there is a much more pronounced magnetic field throughout the plastic region for the low density case without collisions than for the higher densities. Importantly, the turbulent field region at the gold-plastic interface in this case is quite large and quite strong, being made of roughly a half dozen multi-micron scale ‘blobs’ with magnitude on order of 10 times stronger than the filaments in the CH (0.2 vs. 0.02 in simulation units). Watching the evolution of the field in time reveals that these fields both grow in magnitude and move backwards towards the laser focus. This leads to the merging and strengthening of the filaments. As we discuss later these fields insulate the gold from the plastic. These fields impede both the background and high-energy electrons from flowing into the plastic region.

For this low density case, when the Coulomb collisions are switched on, there is hardly any effect, i.e. there are little differences between the two cases. The heat flux plots are quite similar, even down to having two ‘main’ beams plus a few weaker filaments. The B_3 field is also similar, with the one notable difference being that the blobs in the turbulent region at the gold/plastic boundary have perhaps merged even more, creating fewer – and hence larger and stronger – blobs overall. Even fixing the Coulomb logarithm at $\ln(\Lambda) = 50$ doesn’t dramatically change the results, although in the magnetic field plot one can just start to notice an envelope field beginning to surround the filamented region.

3.2.3.4 Discussion of Results

There are a number of different interesting implications from these results. The most fundamental, which can be justified from the results of section 3.2.3.1 alone, is that the increased electron transport to the witness copper as seen in the experiment is primarily due to collimation by the magnetic field in the carbon region, which is greater in the cold foam as




FIGURE 3.13: Forward heat-flux and magnetic field in low density carbon with varied collisions, arranged as in fig. 3.11. It's hard to see much of any change in the behavior between the simulations. Especially in the left two columns, if there is any it is less than the natural variations within an ensemble of simulations with different random seed; with fixed $\ln(\Lambda)$, the q_1 plot still seems similar, although an enveloping B_3 is vaguely apparent in that diagnostic (as not in the other two cases.)

compared to the shocked plasma. The most obvious source for this collimating magnetic field is resistive collimation due to Spitzer resistivity of the return current (see Chapter 1, Section 1.5). In resistive collimation, the curl of the resistive E field leads to the growth of a magnetic field which then guides the forward moving electrons. As resistivity/conductivity in the Spitzer regime is independent of density, we can deduce that the key differences between the two cases are due to differences in thermal profile rather than plasma density.

A naïve conclusion would be that the foam was more collisional because it was initially colder; however as we have shown in section 3.2.3.2 the initial temperature has no bearing on the simulation results whatsoever. As explained there, the reason is that even 40eV is also a negligible temperature compared to those reached due to the heating caused by the high-intensity laser in the gold, the energy from which would then heat the plastic as well. Instead, therefore the difference between the two cases is the relative *heat capacity* of the foam and plasma. For a fully ionized plasma the heat capacity is a linear function of density (for a given ion species,) and so we reach the conclusion that differences in conductivity in the simulations *are* due to density, albeit in a somewhat indirect way.

As has been mentioned multiple times now, the final temperatures reached in the simulations are almost certainly *not* indicative of the actual temperatures of the materials in the experiment – the reason being just about every approximation we have made to make the simulation tractable has the effect of reducing the heat capacity of the target, the most important probably being the greatly reduced plasma density of the gold layer. However, we argue first, that the total absorption of the high-intensity laser as seen in the simulation will not be a strong function of the gold ionization state; and so the absorption fraction measured in the simulation (on order of 15-20%) is a good indication of the absorption fraction from the experiment. Second we argue that the amount of energy therefore deposited – roughly 150J – is more than enough to raise the temperature of the unshocked foam well above 40eV in most reasonable energy partitions between different layers and regions. Therefore we believe the initial temperature of the different materials will not have bearing on the experimental result as well, just as it doesn't in the simulations, and so the main cause of the experimental differences is the different densities (and hence heat capacities). However,

the electron temperature in the plastic will also be substantially less than the temperature seen in the simulations, so that the plastic is more collisional than in the simulations as well. This is the reason we “enhance” the collisions by increasing the $\ln(\Lambda)$.

There is another important difference in the magnetic field between the two cases besides the collimating field in the plastic region, and that is the field which exists at the interface between the gold and the plastic. Looking at the fields shown in figure 3.8, we see that in the case of low density CH there is a large and strong turbulent field on axis in the plastic region just next to the gold foil, and a thinner but still strong field along the gold surface out to the edges of the target; almost no such “surface” fields can be seen in the high density CH target.

The source of this field is two or perhaps three of the other terms besides the collimating term in the generalized Ohm’s law, eq. 1.7. The fields along the Au surface going to the target edge are caused by the $(c/e)\nabla n_e \times \nabla T_e/n_e$ term, as the density gradient (from gold to plastic) and the temperature gradient (from laser focus to box edge) are crossed. The stronger, turbulent region is primarily caused by the $\nabla \times (\mathbf{v} \times \mathbf{B})$ term, which causes any magnetic field generated in the bulk plastic region to be transported back towards the laser focus along with the return current. When comparing the density and B field diagnostics, we see strong field exists even in the rarified regions which is somewhat inconsistent with this analysis and may be due to the $c\nabla n_e \times \nabla |B^2|/(4\pi n_e^2)$ term; however we may also be stretching this Ohm’s law beyond its validity a bit.

The effect of this surface magnetic field will be to impede electron transport from the Au to the plastic region. In the turbulent region behind the laser focus, even the most energetic electrons of 3 MeV will have Larmor radii smaller than the region’s thickness. The field away from the laser focus would allow such an electron to penetrate into the plastic with only partial deflection, though it could impede the flow of thermal electrons or even partially slowed-down $2a_0$ electrons.

Viewing the heat flux and magnetic field results as a function of density and collisionality also reveals several important points. One is that, as a function of density, the plasma

behavior goes from being highly dependent on collisions to almost independent as we go from the high to the low density case; the medium density case is as expected somewhat between the two extremes, though qualitatively more similar to the high rather than to the low density case. Focusing on the high density carbon simulations and looking at the behavior as a function of collisionality we see three distinct behaviors in the three cases investigated. In the collisionless case we see no collective behavior affecting the heat-flux. In the medium collisionality case, we see a local collimation into filaments, which overall gives some average forward collimation as the filaments are generally pointed in the forward direction, especially so for the strongest filaments. However, moving left to right, once the filaments form they travel mostly independently of one another, except for the occasional merger of adjacent filaments; there is no global structure to the collimation. This is in contrast to the behavior when strong collisionality is modeled, where although some filamentary structure is present the more fundamental aspect of the heat-flux is that it has all merged into a single, highly-collimated forward going beam. Interestingly, in the simulation with medium-density carbon and strong-collisionality case, we can see a somewhat ‘in-between’ behavior, as the filaments have coalesced into not one but two separate beams.

We argue that the simulations which most closely match reality are those involving the strong collisionality, because the heat capacity of the gold layer in the simulations is low. In reality the Coulomb logarithm isn’t a free parameter, and so collisionality isn’t something which can be adjusted (outside of having different ion species of course). From that perspective the fixed-logarithm simulations are somewhat artificial. However, as we keep mentioning, it is almost certainly the case that the temperature of the plastic layer is artificially high in the simulations as compared to experiment, because the gold gets far too hot and this heat flows into the plastic. This would have the effect of artificially enhancing the conductivity of the plasma so that the return current creates only small ohmic E-fields; and fixing the Coulomb logarithm at a higher level than would be expected is an ad hoc way of counteracting this. Of course we have no idea what the ‘correct’ collision strength we would need to be consistent with the experiment, and so we would imagine reality to likely be somewhere between the two collisional cases we show here. Even the partial collimation seen in the

lower collisionality simulation could reasonably lead to the transport behavior seen in the experiment, if the target geometry was full-scale. Furthermore, it is very encouraging that enhancing the collisions, and so cooling the plasma, does *not* affect transport in the low density plastic. This means that the relative behavior between the two experiments is not strongly dependent on the details of the background temperature and no collimation occurs for low density while higher plastic density can lead to collimation.

3.3 Summary

We have used the PIC code OSIRIS to investigate a warm-dense matter electron transport experiment. The experiment appeared to show diminished electron transport through a plastic WDM layer compared to a cold and somewhat denser plastic foam. Simulations indicate that the most important physical processes governing the transport behavior are Coulomb collisions and the growth of magnetic fields; and that the most important physical parameter separating the two cases is the density of the transport layer. The initial temperature is found to not be particularly important.

In both cases it is the response of the plasma to the return current which governs the growth of the magnetic field. In the high density plastic, a collimating field region builds up around the return current, due to the resistive electric field cause by Coulomb collisions. In the low density plastic, a turbulent and reflecting field region builds up near the electron source, due to convection of magnetic field back towards the source by the return current.

The simulation most likely over estimates the heating of the target, due to various approximations reducing the heat capacity of the gold layer of the simulated target. The plastic layer heats up from the flux of heat from the gold leading to an artificially temperature in the plastic. This lowers the resistivity of the plastic. To compensate for this we increase the collisionality of the plastic by setting the $\ln(\Lambda)$ to 50. We find that for high foam density this leads to a resistive E-field driven by the return current. This creates a magnetic field from Faraday's law which collimates the forward heat flux carried by the relativistic electrons. For low foam densities the plastic has heated to sufficiently high temperatures that even when

$\ln(\Lambda)$ is set to 50 no collimation is seen.

CHAPTER 4

Macro-Particle Stopping in Particle-in-Cell Simulations

4.1 Introduction

Particle-In-Cell (PIC) simulations [7, 8] are a useful tool for studying nonlinear and kinetic physics such as those found in laser-solid interactions [84, 10], and fast ignition inertial fusion energy [1, 102, 65, 103, 104], which are the focus of this dissertation, as well as relativistic collisionless shocks [105, 106]. Since PIC codes are based on first principles they, along with Vlasov and Vlasov Fokker-Planck codes [107, 108], are a preferred tool to explore physics in regimes where details of the distribution function affect the overall behavior of the system, where there have not been many experiments, or where experimental results are not well understood. In relativistic laser-solid interactions, fast ignition, and relativistic shock studies fully kinetic simulations that resolve the electron dynamics are necessary, and PIC simulations have been used extensively. As discussed earlier in this dissertation, in such problems the flow of relativistic electrons in a background plasma is an important process. However, the physical scale length that needs to be resolved in these systems is the collisionless skin depth and not the Debye length. As the Debye length needs to be resolved in order to avoid numerical grid heating which results from aliasing, the resolution used in such simulations is sometimes much finer or the electron temperature is much higher than necessary. The use of higher order splines, and current smoothing and compensation can eliminate grid heating permitting larger cells and/or lower temperatures to be used [8, 58]. The use of larger cells also permits larger simulations that model the full spatial domains of the problems of interest. In order to carry out such large simulations in multiple dimensions it is common to limit the number of particles per cell, which is generally chosen at a value

that keeps spurious finite size particle collisions and noise to a sufficiently small value.

In this chapter, another effect that needs to be considered in PIC simulations when relativistic electrons are present [109, 110] is described and quantified. This work was originally motivated by studies of fast ignition relevant plasmas [65] and other work in this dissertation. In these simulations a high intensity laser self-consistently generates a distribution of relativistic electrons with densities up to $\sim 10^{22}\text{cm}^{-3}$ moving through a $10^{23}\text{cm}^{-3} - 10^{24}\text{cm}^{-3}$ density background plasma [102, 65, 46]. In a previous paper [65] we observed an anomalous relaxation of the tail in the distribution function of these high-energy electrons in 10^{23}cm^{-3} density plasma; we ascribed this relaxation to the effect of plasmon emission, or wake excitation, in a turbulent plasma. Although plasmon emission is a physical effect, in this chapter we show that its importance on a single particle is enhanced by the use of macro-particles in PIC codes. We show that the stopping power of relativistic electrons due to plasmon emission scales as q^2/m , where q is the charge and m is the mass of the simulated relativistic electrons. In PIC simulations the ratio q/m of the simulated particles is kept consistent with real electrons, but the mass and hence the charge can be many times greater, thus strongly affecting the energy exchange between relativistic electrons and the background plasma. In addition, the stopping power will depend on the dimensionality of the simulation and on the shape of the macro-particles. This artificially high plasmon emission can be controlled by using more particles per cell, larger cells, and higher order particle shapes. This is similar to how collisions and fluctuations are controlled in a PIC code. In this chapter, we derive expressions for the stopping power of relativistic electrons in PIC simulations, and discuss how parameters can be chosen to reduce this effect with an emphasis on two-dimensional (2D) simulations of fast ignition relevant plasmas. These results are also relevant to general laser-solid interaction and relativistic shock simulations.

This chapter is organized as follows. In Section 4.2 the equations for the macro-particle stopping of high energy particles in 2D PIC simulations by energy transfer through wakes are derived. The 1D and 3D cases are also discussed. In Section 4.3 the importance of macro-particle stopping in fast ignition simulations is discussed and the stopping of single particles in a nominal fast ignition plasma is analyzed. Section 4.4 addresses the relaxation of high-

energy tails of the electron distribution function due to macro-particle stopping. In Section 4.5 we turn to full-scale fast ignition simulations and analyze how varying the number of particles per skin depth in isolated target simulations changes the results. In section 4.6 a particle splitting technique we have developed is presented which reduces the effect of stopping by reducing the charge on each high energy electron. Finally, in Section 4.7 we summarize our results.

4.2 Theory

A charged particle moving faster than the thermal velocity in a plasma forms plasma wave wakes. The electric field of the wake at the location of the particle will do work on the particle and slow it down. To find the energy loss due to these wakes, we will consider a highly relativistic point electron moving through a cold fluid background plasma, and will solve for the electric field of the wake, following the methods used in the plasma-based accelerator community [111, 112]. We present results for the wake and electric field on the particle for 1D, 2D, and 3D simulations. We then consider the effect of finite sized particles on the stopping, with an emphasis on 2D as it is presently the most common situation.

We start by considering relativistic particles with charge ρ_b moving near the speed of light (c) in a cold fluid plasma, following the approach in Ref. [111]. The linearized fluid equations for the background plasma are:

$$\frac{\partial \vec{v}_1}{\partial t} = -\frac{e}{m_e} \vec{E}, \quad (4.1)$$

$$\frac{\partial n_1}{\partial t} = -n_0 \vec{\nabla} \cdot \vec{v}_1, \quad (4.2)$$

$$\vec{\nabla} \cdot \vec{E} = -4\pi e n_1 + 4\pi q n_b, \quad (4.3)$$

where e is the charge of an electron, m_e is its mass, v is the electron fluid velocity, E is the electric field, n is the plasma density, the subscripts 0 and 1 indicate zeroth and first order quantities respectively, and the subscript b corresponds to quantities for the relativistic

particle species. We next consider the response for a single beam particle moving in the z direction with a speed v_b , in which case

$$qn_b = q\delta(\vec{r} - \vec{v}_b t) = q\delta(z - v_b t)\delta(\vec{r}_\perp) \quad (4.4)$$

where $\delta()$ is the Dirac delta function, q is the magnitude of the charge on the test particle, $\vec{r}$ the position vector, and $\vec{r}_\perp$ is the part of the position vector perpendicular to the direction of motion. The response to this point particle can also be viewed as the Green's function response of the electric field when one is calculating the response to a continuous beam distribution. We will also use Faraday's law

$$-\vec{\nabla} \times \vec{E} = \frac{1}{c} \frac{\partial \vec{B}}{\partial t}, \quad (4.5)$$

and Ampere's law where we have substituted for the current

$$\vec{\nabla} \times \vec{B} = -\frac{4\pi en_0}{c} \vec{v}_1 + \frac{4\pi qn_b}{c} \vec{v}_b + \frac{1}{c} \frac{\partial \vec{E}}{\partial t}. \quad (4.6)$$

We combine these equations to get

$$\frac{\partial^2 n_1}{\partial t^2} + \omega_p^2 n_1 = \omega_p^2 \frac{q}{e} n_b = \frac{\omega_p^2 q}{v_b e} \delta(\vec{r}_\perp) \delta(t - \frac{z}{v_b}). \quad (4.7)$$

where $\omega_p = \sqrt{4\pi n_0 e^2 / m_e}$ is the unperturbed plasma frequency.

Due to causality the density perturbation in front of the relativistically moving test particle must vanish, so the density response is

$$n_1 = \frac{\omega_p q}{v_b e} \delta(\vec{r}_\perp) \eta(t - \frac{z}{v_b}) \sin(\omega_p(t - \frac{z}{v_b})) \quad (4.8)$$

where $\eta()$ is the Heaviside step function. To find the electric field, we use equations (4.1), (4.5), and (4.6) to derive a wave equation for the electric field,

$$-\frac{1}{c^2} \frac{\partial^2 \vec{E}}{\partial t^2} + \nabla^2 \vec{E} - \vec{\nabla} \vec{\nabla} \cdot \vec{E} = k_p^2 \vec{E} + \frac{4\pi}{c^2} \frac{\partial \vec{v}_b q n_b}{\partial t}. \quad (4.9)$$

where $k_p = \frac{\omega_p}{c}$ is the plasma wave number.

Next, we substitute from Gauss's law (4.3) and assume the test particle's velocity remains very close to c , i.e. $\vec{v}_b = c\hat{z}$, where $\hat{z}$ is the unit vector in the z direction. We concentrate on

the component of the electric field in the $\hat{z}$ direction, and use the fact that both E_z and ρ_b are therefore functions of $(z - ct)$ such that $\frac{\partial}{\partial t} + c\frac{\partial}{\partial z} = 0$, to obtain

$$(\nabla_{\perp}^2 - k_p^2)E_z = -4\pi k_p q \frac{\partial}{\partial z} [\delta(r_{\perp}) \eta(t - \frac{z}{c}) \sin(\omega_p(t - \frac{z}{c}))] \quad (4.10)$$

The solution for E_z can then be written as

$$E_z = G_R(r_{\perp}) 4\pi k_p^2 q \eta(t - \frac{z}{c}) \cos(\omega_p(t - \frac{z}{c})) \quad (4.11)$$

where

$$(\nabla_{\perp}^2 - k_p^2)G_R(r_{\perp}) = \delta(r_{\perp}). \quad (4.12)$$

The solutions to Eq. (4.12) depend on the dimensionality of the problem. In 1D, the $\delta(r_{\perp})$ in Eq. (4.4) should be replaced by 1, $\nabla_{\perp}^2$ vanishes, and q is in units of charge per unit area, so that $G_{R1D} = -1/k_p^2$. This yields an expression for the electric field

$$E_{z1D}(z) = -4\pi \bar{q} \eta(t - \frac{z}{c}) \cos(\omega_p(t - \frac{z}{c})) \quad (4.13)$$

where $\bar{q}$ is the charge per unit area.

In 3D or 2D cylindrical, the solution to Eq. (4.12) can be shown to be the Modified Bessel Function of the Second Kind [113]

$$G_{R3D}(r_{\perp}) = -\frac{1}{2\pi} K_0(k_p |r_{\perp}|) \quad (4.14)$$

which yields the Green's function for the electric field behind a test charge

$$E_{z3D}(r_{\perp}, z) = -2q k_p^2 K_0(k_p r_{\perp}) \eta(t - \frac{z}{c}) \cos(\omega_p(t - \frac{z}{c})), \quad (4.15)$$

The stopping power on a charge is the Green's function evaluated at the particle. For point particles the value of the step function is $1/2$ at the position of the particle. From Eq. (4.15), in 3D the electric field diverges logarithmically at the origin, which violates the linearity condition. The stopping power would also diverge logarithmically. Therefore, if this response is viewed as the electric field from a real point electron then the result cannot be valid near the origin. However, when viewed from the stand point of a finite size particle (or beam), what matters is the solution for E_z obtained by integrating the Green's function over the

particle's shape. As long as this electric field is finite then the use of linear theory can still be valid. We next carry out this integral in detail for the 2D case and comment on the 1D and 3D cases at the end of the section.

In 2D (slab) geometry, where $\nabla_{\perp}^2 \rightarrow \frac{d^2}{dr_{\perp}^2}$ and $\delta(r_{\perp}) = \delta(r_{\perp})$ it is readily verified [112] that the solution to Eq. (4.12) is given by setting $G_{R2D}(r_{\perp}) = -\frac{1}{2k_p}e^{k_p r_{\perp}}$ for $r_{\perp} < 0$ and $G_{R2D}(r_{\perp}) = -\frac{1}{2k_p}e^{-k_p r_{\perp}}$ for $r_{\perp} > 0$. Thus the electric field is

$$E_{z2D}(r_{\perp}, z) = -2\pi e^{-k_p |r_{\perp}|} k_p \bar{q} \eta(t - \frac{z}{c}) \cos(\omega_p(t - \frac{z}{c})), \quad (4.16)$$

where $\bar{q}$ is the charge per unit length. Unlike the 3D case, this electric field remains finite at the origin allowing direct analysis without specifying a particle shape. Using $1/2$ for the value of $\eta(0)$, the energy loss for a point (line of charge in 2D) particle is given by

$$\frac{d\epsilon_{2D}}{dt} = -\pi \omega_p \bar{q}^2 \quad (4.17)$$

where ϵ is the energy of the test particle. Dividing by $\omega_p \bar{m} c^2$ we get

$$\frac{d\gamma_{2D}}{d\omega_p t} = -\pi \frac{\bar{q}^2}{\bar{m} c^2}. \quad (4.18)$$

where γ is the relativistic Lorentz factor of the test charge. Equation (4.18) shows that the stopping power (distance) for a relativistic particle in a cold plasma is (inversely) proportional to the $\bar{q}$ on the particle holding the $\bar{q}/\bar{m}$ ratio fixed. To derive an equation for PIC simulations we use the fact that $\bar{q}/e = n_0 \Delta^2 / N$ where Δ is the cell size (we assume equal cell sizes along the different dimensions) and N is the number of simulation macro-particles per cell, and obtain

$$\frac{d\gamma_{2D}}{d\omega_p t} = -\frac{1}{4} \frac{\omega_p^2}{c^2} \frac{\Delta^2}{N} C. \quad (4.19)$$

Equation (4.19) demonstrates that for large cell sizes N must be increased to reduce the macro-particle stopping to a desired level. The factor C is a term that accounts for finite size particle effects which will depend on the details of the particle shape and smoothing used in the simulation.

Next we consider finite size particle effects and the factor C . These effects will reduce the stopping power of the test particle. The use of finite size particles reduces both the electric

field made by the test charge and the force from this electric field on the particle. We first discuss the reduction of the electric field. The effect of the finite size of the particle on the electric field is found from a convolution of the point particle electric field with the shape factor as a function of grid size

$$\bar{E}_z(r_\perp, \xi) = \int_{-\infty}^{\infty} dr'_\perp \int_{-\infty}^{\infty} d\xi' E_z(r_\perp - r'_\perp, \xi - \xi') S(r'_\perp, \xi'), \quad (4.20)$$

with $\xi = (z - ct)$. For linear splines (area weighting) in one dimension the shape factor is given by

$$S_1(x) = \frac{1}{\Delta} \left(\left(1 - \frac{x}{\Delta}\right) \eta(x) \eta(\Delta - x) + \left(1 + \frac{x}{\Delta}\right) \eta(-x) \eta(\Delta + x) \right), \quad (4.21)$$

where in Eq. 4.20, x would be one of $(r'_\perp, \xi')$ and Δ is the cell size in a specified direction; the full shape function is the product of the one dimensional shape functions for each dimension. As noted above we consider square cells in 2D ($\Delta_{r_\perp} = \Delta_\xi = \Delta$) but the arguments presented can be easily extended to rectangular cells. The Fourier transform of this shape function, which is an order 1 spline, is a sinc function squared. Higher order spline shapes are defined as the convolution of the previous order with the zeroth order shape (nearest grid interpolation),

$$S_n(x) = \int_{-\infty}^{\infty} dx' S_0(x - x') S_{n-1}(x'). \quad (4.22)$$

Therefore, for splines of order n the Fourier transform of the shape function is given by

$$S(k)_n = \left(\frac{\sin(\frac{k\Delta}{2})}{\frac{k\Delta}{2}} \right)^{n+1}. \quad (4.23)$$

The shape factor is dependent on the cell size and on the interpolation order. In addition, digital filters are also commonly applied. A filter of particular interest used for current smoothing is a triangular (1-2-1) filter applied multiple times in each direction. A compensator [8] is often used to flatten the profile in k -space for low k . These filters take the k -space form

$$F(k)_n = \cos^{2n} \left(\frac{k\Delta}{2} \right) \quad (4.24)$$

where here n is the number passes of the triangular filter. When the compensator is included the effective filter takes on a more complex expression, but there is no k dependence in the




FIGURE 4.1: *Theoretical curves for the magnitude of the z component of the on axis electric field of a relativistic particle moving in the z direction. a) The strength of the electric field is dependent on cell size with larger cell sizes producing weaker fields. b) The strength of the electric field is also dependent on the interpolation order with higher order interpolation schemes producing weaker fields.*

denominator so smoothing (including compensation) is never as effective as using higher-order splines at reducing aliased noise (the source of particle heating in PIC codes).

For illustrative purposes, we examine the electric field generated by different particle shapes where the charge on the particle is kept fixed. Figure 4.1 shows the magnitude of the axial (z) component of the electric field produced on axis ($r_\perp = 0$) as a function of z by a 2D relativistic particle of finite size (square) in a plasma according to equation (4.20). In figure 4.1a the particle shape is a second order spline, i.e., $n = 2$ in Eq. (4.23), and the cell size is varied from a point particle, to $0.5c/\omega_p$, to $1.0c/\omega_p$, and to $2.0c/\omega_p$ (blue, red, green

and magenta curves respectively). These curves confirm that increasing the cell size for a fixed charge reduces the electric field strength and smooths it out over the location of the particle. In Figure 4.1b, we show that increasing the interpolation order while keeping the cell size and charge on the particle fixed also decreases the amplitude of the electric field. We plot the electric field for linear (blue), quadratic (magenta), cubic (green), and quartic (red) interpolation while holding cell size fixed at $2.0c/\omega_p$. This shows that each higher order spline corresponds to a ‘wider’ particle. Varying the cell size or varying the particle order both effectively change the size of the particle, however, the use of larger cells also changes the accuracy of the field solver for how light waves propagating in vacuum. We also note that although the macro-particle stopping effects are reduced as the cell size is increased to values larger than a skin depth (and Debye length), other issues such as significant modifications of the dispersion properties of both plasma and electromagnetic waves will become important.

Next, we turn our attention from the electric field of the wake to the stopping power of the relativistic particles. The stopping power is obtained from an integral of the convoluted electric field $\bar{E}_z$ in Eq. (4.20) with the shape function

$$\begin{aligned} F_{stopping} &= q \int_{-\infty}^{\infty} d\vec{r}_{\perp}' \int_{-\infty}^{\infty} d\xi' \bar{E}_z(\vec{r}_{\perp}', \xi') S(\vec{r}_{\perp}', \xi') \\ &= q \int_{-\infty}^{\infty} d\vec{r}_{\perp}' \int_{-\infty}^{\infty} d\xi' \int_{-\infty}^{\infty} d\vec{r}_{\perp}'' \int_{-\infty}^{\infty} d\xi'' E_z(\vec{r}_{\perp}' - \vec{r}_{\perp}'', \xi' - \xi'') S(\vec{r}_{\perp}', \xi') S(\vec{r}_{\perp}'', \xi'') \end{aligned} \quad (4.25)$$

It is worth noting that total force on the particle can also be obtained by integrating the shape function and electric field in k space,

$$F_{stopping} = q \int_{-\infty}^{\infty} d\vec{k} S^2(\vec{k}) E_z(\vec{k}). \quad (4.26)$$

Through Eq. (4.21) it can be seen how the use of progressively higher order splines (and smoothing) can reduce the stopping force. The shape functions can be different in each direction, but it is generally a separable function in the each coordinate in which only the cell size varies. For separable shape functions, it is easy to extend eqs. 4.25 and 4.26 to multi-dimensions. For simplicity, in what follows we assume the cells are squares in 2D.

It is difficult to obtain analytical expressions for $F_{stopping}$ for the particle shapes used in PIC codes, therefore we simply plot numerical solutions to Eq. (4.25). Figure 4.2 shows the




FIGURE 4.2: Reduction in stopping force due to finite-sized particles (“ C ” in eq. (4.19).) ‘Smoothing’ is used in finite-difference codes like OSIRIS, and here refers to a 4-pass (1,2,1) filter followed by a (-5,14,-5) compensator [8]. ‘Filtering’ is done directly in k -space and therefore is used in spectral codes like PARSEC; here the filter is $e^{-\frac{k^2 \Delta^2}{2}}$, where Δ is the cell size (applied to both the particles and the fields). From the figure it can be seen that smoothing is more effective at reducing stopping than higher-order particle shapes, and filtering is more effective than smoothing.

effect of varying the spline order and current smoothing on the stopping force. We plot the force on the particle for linear and cubic splines, and the same using 4-pass smoothing (which we often use in OSIRIS) and k -filtering (used in PARSEC); it is normalized to the stopping for a point particle, and therefore this curve is equal to the factor C for two dimensions, C_2 , in Eq. (4.19). As stated earlier, both smoothing and higher-order interpolation schemes reduce the macro-particle stopping. However, for fixed charge per particle, the reduction is only significant for large cell sizes ($\Delta \geq 1/k_p$), where important skin-depth physics can be missed and where significant numerical dispersion effects can be present [8]. In essence by increasing the cell size, one is smoothing out the wake field, whose scale length is on order of the skin depth. For example, for a cell size of $1/2 k_p^{-1}$, without current smoothing, the stopping power is 77% and 68% of the point-particle limit for linear and cubic interpolation respectively. If the cell size is increased to $3/2 k_p^{-1}$, then these values become 37% and 21% respectively. Including current smoothing further reduces the value of C to 0.09 in the linear case and 0.05 in the cubic one. We note that there are numerical limits to both increasing cell

size and increasing interpolation order for the purpose of reducing macro-particle stopping. For example, the use of large cells can lead to plasma waves and light waves having negative group velocity. In finite-difference PIC codes where numerical errors cause the phase velocity of electromagnetic waves on the grid to be less than the speed of light, relativistic particles can excite numerical Cerenkov noise which is increased as the time step is reduced below the electromagnetic Courant condition for stability [8]. In addition, since the hard stability limit for the time step in a plasma is $2/\omega_p$ [8] this introduces a practical limit on the cell size of $2\sqrt{2}c/\omega_p$ when dealing with relativistic particles. As seen in equation (4.19) the stopping power can be reduced by also increasing the number of particles per cell (or effectively decreasing the particle charge). This is the best method with respect to not modifying the physics, but it is obviously the most computationally costly method.

We close this section with a brief discussion on macro-particle stopping in 1D and 3D. The electric field for an infinitesimally thin sheet of charge was given in Eq. (4.13). We note that it could also be obtained from integrating the 3D result, Eq. (4.15), over a beam with a charge density that is uniform across the beam in the transverse direction. Interestingly the 1D result is independent of the plasma density, however, the result is only valid if the electric field is still sufficiently small for linear theory to be valid. This is true if the normalized field value is small, $\frac{eE_z}{mc\omega_p} \ll 1$ [112, 111]. For a macro-particle one needs to integrate Eq. (4.13) over the shape of the particle. The energy loss for a point (sheet of charge in 1D) particle is given by

$$\frac{d\epsilon_{1D}}{dt} = -2\pi\bar{q}^2c, \quad (4.27)$$

which can be written for the Lorentz factor of the test charge as

$$\frac{d\gamma_{1D}}{d\omega_pt} = -2\pi\frac{\bar{q}^2}{\bar{m}\omega_pc}. \quad (4.28)$$

Taking into account the fact that $\bar{q}/e = n_0\Delta/N$ in 1D, we obtain the energy loss in 1D PIC simulations

$$\frac{d\gamma_{1D}}{d\omega_pt} = -\frac{1}{2}\frac{\omega_p}{c}\frac{\Delta}{N}C_1 \quad (4.29)$$

where C_1 is the effect of the particle shape in 1D.

In 3D the Green's function logarithmically diverges as $|\vec{r}_\perp| \equiv r$ approaches zero. However, for a finite size particle the divergence disappears. A detailed analysis is beyond the scope of this work. However, in [114] the result for flat top and Gaussian particle (beam) shapes in the transverse coordinate and Gaussian shapes in the longitudinal directions were given. Here we summarize the results for Gaussian shaped macro-particles, for which the charge is given by $\frac{q}{(2\pi)^{3/2}} e^{-(z-ct)^2/2\sigma_z^2} e^{-r^2/2\sigma_r^2}$, where σ_z and σ_r are the particle size in the longitudinal and transverse directions respectively. For such a charge the amplitude of the wakefield is given in [114] as

$$E_{z3D} = -qk_p^2 e^{-k_p^2 \sigma_z^2/2} e^{k_p^2 \sigma_r^2/2} \Gamma\left(0, \frac{k_p^2 \sigma_r^2}{2}\right), \quad (4.30)$$

where $\Gamma(\alpha, \beta) = \int_\beta^\infty s^{\alpha-1} e^{-s} ds$ is the incomplete gamma function. For symmetric macro-particles, $\sigma_r = \sigma_z = \sigma$, we have simply

$$E_{z3D} = -qk_p^2 \Gamma\left(0, \frac{k_p^2 \sigma^2}{2}\right). \quad (4.31)$$

We note that in the limit that $k_p \sigma$ approaches zero, $\Gamma(0, \frac{k_p^2 \sigma^2}{2}) \approx \ln(1.12/k_p \sigma)$, and the electric field on axis reduces to

$$E_{z3D} = -qk_p^2 \ln\left(\frac{1.12}{k_p \sigma}\right). \quad (4.32)$$

This is the peak amplitude of the wake. To derive how the electric field is distributed over the particle requires carrying out the full integrals over the particle which is beyond the scope of this work. However, this expression is useful for estimating the macro-particle stopping and we will use it in what follows.

Following the same procedure as in the 1D and 2D cases, the energy loss in 3D is then described by

$$\frac{d\epsilon_{3D}}{dt} = -q^2 \frac{\omega_p^2}{c} \ln\left(\frac{1.12}{k_p \sigma}\right), \quad (4.33)$$

or, for the Lorentz factor of the test charge, as

$$\frac{d\gamma_{3D}}{d\omega_p t} = -\frac{q^2 \omega_p}{m c^3} \ln\left(\frac{1.12}{k_p \sigma}\right). \quad (4.34)$$

If we assume cubic cell shapes with the same cell size, Δ in each direction then $q/e = n_0 \Delta^3/N$ and we use the cell size for the particle size, σ , then we can estimate the macro-particle

stopping in 3D, which is given by

$$\frac{d\gamma_{3D}}{d\omega_p t} = -\frac{1}{4\pi} \frac{\omega_p^3}{c^3} \frac{\Delta^3}{N} \ln \left(\frac{1.12}{k_p \Delta} \right) C_3. \quad (4.35)$$

where C_3 is a factor that accounts for additional macro-particle effects. We note that for a fixed number of particles per cell, N , and a cell size $\Delta \leq c/\omega_p$ (neglecting the exact details of the particle shape), the macro-particle stopping decreases with increasing dimensionality. Therefore, it requires a larger number of particles per cell in 1D and 2D to control this effect when compared with 3D. We note that the formalism that has been presented can be used to carry out an exact and detailed analysis of the 3D case as we did for the 2D case.

4.3 Single Particle Stopping in Fast Ignition Simulations

We can now make predictions for the stopping of relativistic macro-particles in background plasmas. Here we focus on the stopping distance in two dimensional fast ignition relevant simulations. To motivate the importance of macro-particle stopping for fast ignition studies, let us take as example the parameters used in integrated PIC simulations of fast ignition [65, 102, 103] aimed at modeling both laser absorption and transport of fast electrons in a background plasma. We compute the stopping power for the simulations with different resolutions and numbers of particles per cell. In these simulations an intense laser with a $1 \mu\text{m}$ wavelength interacts with a plasma with a nominal density around $100n_c$ (where n_c is the critical plasma density for the laser). The number of macro-particles per skin depth squared (PPSD²) in these simulations varied between 6 and 24, the cell sizes varied from 0.5 to $2 k_p^{-1}$, and the particle order varied from linear to cubic. Despite these differences, the predicted energy loss rate was similar and it only varied in the range $0.12 - 0.34 \text{ MeV}/\mu\text{m}$. The simulations used different plasma lengths so the single particle energy loss across the simulation box varied between 4 and 18 MeV . For example, in the paper by Tonge et al.[65] the cell size was $0.5 c/\omega_p$, with 4 particles per cell ($N = 4$) giving 16 PPSD^2 , and quadratic splines and 5-pass compensated smoothing were used, leading to a stopping force of $0.011 m_e c \omega_p$, or an energy loss rate of $0.337 \text{ MeV}/\mu\text{m}$ at a density of $100n_c$. In each of these cases the energy loss of a macro-particle due to enhanced stopping over the simulation box

size is comparable or larger than the typical fast-electron energies relevant for fast ignition. Therefore the enhanced stopping of macro-particles could greatly affect the results in each case.

We verify these predictions by running simulations of single particle stopping with the finite difference PIC code OSIRIS [2, 115, 116]. The simulation setup is as follows. The cell size is $0.5 c/\omega_p$, which in OSIRIS is also the particle size. For the background plasma, both ions and electrons are included with a mass ratio of 3672 (Deuterium plasma); the particle count is 64 particles per cell, and the temperature is 100eV. For the beam, a single electron is accelerated to a forward momentum of $100m_e c$, over a time period of $300/\omega_p$ to minimize radiation from the particle acceleration (without allowing the plasma response to affect the particle), and given a weight ($\bar{q}$) equivalent to 16 particles per skin depth squared. The box size is $25 c/\omega_p$ in the transverse dimension, and long enough in the longitudinal direction that the particle does not have time to cross the box; boundary conditions are periodic in all directions. Linear splines are used, and we use smoothing and compensation for the current. In order to study the wake driven by a single electron it is necessary to remove the effects on the field caused by the thermal motion of the background electrons, as otherwise these would dominate. To do this we use a subtraction technique [117, 118] wherein a second simulation is run with an identically initialized background plasma but without the beam particle, and the difference of the fields is taken to show the effects of the beam alone.

Figure 4.3a shows the E_1 obtained using the subtraction technique; the wake can clearly be seen. In figure 4.3b we show a line-out of the wake field on axis from 4.3a (red), zoomed in around the position of the particle (now taken to be the origin), along with the theoretical prediction for the wake, including the effect of finite sampling, finite size particles, current smoothing and compensation (blue). As can be seen the results are in good agreement, although of somewhat reduced magnitude.

We also compared the energy loss observed in these simulations to the theory presented in Section 4.2. Looking at the energy loss over a period of $500/\omega_p$, Eq. (4.19) predicts a $\Delta\gamma$ of 7.8125; the factor C is calculated to be 0.70 for particles of this shape, giving an adjusted $\Delta\gamma$ of 5.4616. In our simulation we see a $\Delta\gamma$ of 5.2448, 96% as much as is predicted by




FIGURE 4.3: a) Wake formed in the E_1 field by a relativistic particle traveling along the x_1 direction in a 2D PIC simulation. This wake is predicted by wakefield theory. b) A comparison of the lineout along x_1 of the E_1 field near the relativistic particle predicted by theory (blue) and from the PIC simulation (red). The curves are in close agreement, though the simulation field is slightly weaker.

theory; this is consistent with the slightly weaker field seen in Fig. 4.3b.

We postulated that the deviation from theory might be due to the fact that the background plasma also is composed of finite size particles, with a size and shape equivalent to the beam particle. This would also have the effect of smoothing out the wake and giving a weaker field. To test this hypothesis, we reduced the cell and particle size but kept other parameters the same (reducing the number of particles per cell to keep the total particle count constant). Unfortunately using a standard Yee [77] field solver, numerical Cerenkov radiation becomes a dominating factor in the energy loss of the particle as the cell size becomes smaller. For a cell size of $1/16k_p$, we find numerical Cerenkov is already contributing more to the stopping than the wakefield. We experimented with various field solvers in OSIRIS to alleviate this issue. We have found solvers based on the work of Pukhov [119] and Lehe [120] to both be effective at suppressing the numerical Cerenkov for particles traveling along the axis (we have not investigated numerical dispersion issues in other directions.) In addition, since the publication of the paper on which this chapter is based, we have developed a field solver for OSIRIS which eliminates this effect [121, 122, 123, 124, 125]. For the same parameters as the previous paragraph but a cell size of $1/16k_p$ the factor C in eq. (4.19) is calculated to be 0.96, which gives a predicted $\Delta\gamma$ of 7.515. In OSIRIS using a Lehe solver, we observe a $\Delta\gamma$ of 7.336, which is 98% of theory. Therefore we conclude that the wakes are also additionally reduced by the use of finite sized particles in the plasma.

We have also performed a series of test particle simulations with PARSEC [126] to examine the effect of background temperature and particle energy. Recall that the theory assumes a cold fluid plasma and that the particle is highly relativistic. These 2D PIC simulations use the same simulation box and numerical parameters of the previous simulations, but we varied the energy of the relativistic electron (> 1 MeV) and of the background plasma (100 eV - 1MeV). For background temperatures up to 10keV and all fast electron energies examined, the energy loss is relatively insensitive to either parameter, and is $\sim 14\%$ less than the theoretical value. For background temperatures > 100 keV, the velocity of background particles becomes relativistic, the assumption of a cold background plasma does not hold and the simulation results deviate more strongly.

We have tested $N = 4, 16, 64$, and 256 for the background plasma and the differences to the stopping of a single particle were negligible for background temperatures below 10keV . We therefore conclude that equation (4.19) is useful for all initial energies above 1MeV and for background temperatures up to $\sim 10\text{keV}$.

4.4 Relaxation of High-Energy Tails

We now investigate the macro-particle stopping for a distribution of fast electrons with a high-energy tail. These simulations are run with PARSEC and model a spatially uniform “cold” (1keV) background plasma with a relativistic tail distribution with a slope temperature of 2.5MeV and with 4% of the particles in the tail. The simulations have periodic boundaries, with a box size of 256 cells by 256 cells, a cell size of $0.5c/\omega_p$, and $N = 16$ or 64 PPSD². Second order particle interpolation is used with Gaussian shaped particles with particle sizes equal to the grid size. In Figure 4.4 the blue curve shows the initial energy distribution, with the red, green and yellow curves showing the relaxation of the distribution with time, 92.75fs apart when scaled to a background density of 10^{23}cm^{-3} . The curves show that the high-energy particles uniformly loose energy as predicted by theory. The predicted energy loss is $0.09\text{MeV}/\mu\text{m}$ ($0.027\text{MeV}/\text{fs}$) and the simulations show an energy loss of $0.06\text{MeV}/\mu\text{m}$, which is less than both theory and single particle simulations.

In single particle simulations the energy loss is 86% of theory and insensitive to background particle count; in simulations with a high-energy tail the discrepancy with theory is dependent on the PPSD² because this affects the distribution of particles in the tail. For simulations where we vary the PPSD² but keep other parameters constant, we found energy losses of 71% , 66% , and 53% of the predicted value for 16 , 64 , and 256 PPSD² respectively. This effect is likely due to the more even distribution of particles in the tail across cell boundaries as N (not necessarily PPSD²) increases. For the particles in the tail the number of particles per cell varied from $N = 0.16$, 0.64 , and 2.56 respectively (recall that 4% of the particles are in the tail). As the tail becomes more uniformly distributed in space the discrete effect of macro-particles is decreased. In simulations relevant to fast ignition the electrons




FIGURE 4.4: *Relaxation of high-energy tails in a periodic simulation of a cold background and a hot electron tail. The blue curve shows the initial energy distribution, with the red, green and yellow curves showing the relaxation of the distribution with time, 92.75 fs apart when scaled to a density of $100n_c$. The curves show that the high-energy particles uniformly lose energy as predicted by theory.*

are accelerated by the laser in bunches [104] separated by half a laser wavelength which is comparable to a plasma wake wavelength. Therefore, the particles in these bunches will not interact through individual wakes but rather through collective wakes. If the bunch duration is short compared to a skin depth the wakes will collectively add leading to enhanced stopping while if the particles are bunched on scales much longer than the skin depth then the stopping power will be comparable to the single particle stopping power.

4.5 Integrated Fast Ignition Simulations

We will now revisit the integrated simulations using isolated fast ignition targets from Tonge et al. [65]. In these simulations a $50\mu\text{m}$ radius $100n_c$ target with a $20\mu\text{m}$ core region is illuminated with an $I = 8 \times 10^{20}\text{W}/\text{cm}^2$, $\lambda = 1\mu\text{m}$ laser with a $20\mu\text{m}$ spot size; detailed parameters are given in Tonge et al.[65]. Here we reproduce these simulations using cubic splines (quadratic splines were used in the original simulations) and varying the PPSD² while keeping all other parameters fixed. By increasing the PPSD² from 16 (original value) to 100 we decrease the stopping by 6.25 times and we observe a significant change in the dynamics of the simulation due to macro-particle stopping. The stopping decreases the energy of hot electrons while increasing the heating of the background plasma between the laser-plasma interface and the target core. Figure 4.5 shows the momentum distribution function in the isolated target, 835 fs after the laser strikes, at the longitudinal (x_1) position $2\mu\text{m}$ behind the laser-plasma interface and integrated along the transverse direction (x_2) for the 16 PPSD² or 100 PPSD² cases. The blue curve (16 PPSD²) is clearly wider than the red curve (100 PPSD²) indicating the plasma is hotter in the interaction region. Panel b) shows the power absorbed in the target core as a function of total laser power. From the initial heating at 400 fs up to 600 fs the two simulations have similar power delivered to the core, although the red curve has a bump at 450 fs which is due to a higher level of refluxing off the back of the target early in the high particle count simulation. After 600 fs the curves diverge, with the lower PPSD simulation showing a larger heating of the core; this is due to both energy of the electrons reaching the core being lower and hence the stopping of fast electrons in the core

being enhanced, and to the increased heating of the background plasma in front of the core that causes heated background electrons to travel into the core and redeposit their absorbed energy. This overall effect was to increase the power delivered to the core by 25% at 1.3 ps. In the lower absorption case (higher PPSD case) the laser is still capable of delivering 10% of its power to the core.

Let us now compare the magnitude of the single particle stopping effect to that seen by a real particle (although these are 2D simulations, the 3D result is equivalent up to a factor of $\mathcal{O}(1)$.) For a $1\mu\text{m}$ laser, $100n_c = 10^{23} \text{ cm}^{-3}$, the cell size of $0.05c/\omega_0$ is $\sim 0.008\mu\text{m}$, giving 5×10^4 electrons in a cell volume. For $N = 25$, and taking into consideration finite particle size effects, the stopping is ~ 1300 times stronger in the simulation than would be seen for a physical electron. To make the real and simulated macro-particle stopping powers numerically equivalent at $n = 100n_c$ and retain the skin depth physics (i.e. cell size $\leq c/\omega_p$) we would have to increase the particle count to $N > 500$ and use current smoothing and cubic splines. This is impractical for multi-dimensional PIC studies of relevant fast ignition scales. However, it is not important to make the macro-particle stopping power comparable to the real stopping power, but rather to make it small enough over the plasma size so that this enhanced stopping does not change the results. On the other hand, if one is modeling the heating of the core at very high densities, then since the stopping power of the electrons is the physics process that is most important it needs to be quantitatively correct. This is similar to the effect of the collisions between macro particles. If one is modeling a region of plasma for which collisions are not important it is not important that the collision frequency is correct, rather it must be kept small enough that the growth rates and dispersion properties of collisionless processes are not modified.

4.6 Particle Splitting

We now discuss a scheme that can potentially overcome this effect of macro-particle stopping in PIC simulations associated with laser-solid interactions, fast-ignition, and relativistic shocks. We have shown that reducing the charge on the relativistic, i.e., ‘beam’, electrons is




FIGURE 4.5: Panel a) shows the distribution function for forward momentum (p_1) $3 \mu\text{m}$ in front of the target core at 835 fs for 16PPSD² and 100PPSD² simulations respectively. These show that macro-particle stopping results in less heating of the plasma. Panel b) compares the power delivered to the core for the 16PPSD² and 100PPSD² simulations. These show that greater macro-particle stopping also results in lower energy electrons carrying energy to the target core.

an effective method to alleviate the enhanced stopping of relativistic particles due to their wakes. However, in the majority of the simulations of interest we want to study the way particles are accelerated, e.g. by the laser in fast ignition scenarios or by the shock in astrophysical plasmas, and we do not know beforehand which particles will be accelerated (fast) and which will be background plasma particles. Thus, in a standard PIC simulation reducing the charge of a fast particle requires reducing the initial charge of all simulation particles, leading to a prohibitive particle count.

In many of these scenarios, the super-thermal particles which are stopped by macro-particle effects are only a small percentage of the total particle count - although at the same time these particles can have a profound effect on plasma dynamics. Under these conditions an algorithm that splits the relativistic electrons into macro particle electrons with less charge might be effective. We have implemented and experimented with a splitting algorithm that works as follows. For a given simulation we define the minimum particle energy above which macro-particle stopping will become important and the maximum particle charge that high-energy particles should have for accurate description of their stopping based on the theory presented in Section 4.2. In the simplest algorithm these can be fixed values, for instance associated with the initial laser and plasma conditions. It would also be possible to make them dynamic values that depend for instance on the local background plasma conditions that vary due to density gradients or plasma steepening. The simulation is initialized in the standard way, however, after each n time steps (in the future the value of n could change statically or dynamically as the simulation progresses), we calculate the energy and charge of each particle and compare it to the minimum energy and maximum charge defined. If they are both larger, then we reduce its charge and mass by a factor of two and then duplicate the particle. This process is repeated until the particle charge is smaller than the defined value. Thus, we effectively increase the particle count for the high-energy part of the distribution function. It is important to note that in most cases, because of the wide difference between the background thermal and ‘fast’ particle energies, the efficiency of the splitting is relatively insensitive to the exact energies chosen for the splitting condition. Also, since the number of fast particles is significantly smaller than the number of background particles,




FIGURE 4.6: (Color) Particle count vs. the logarithm of the kinetic energy (red), and the same data weighting each particle by the heatflux it carries in the x_1 direction (blue.) Arbitrary units, both independently normalized to a peak of 1.

the computational overhead of splitting is not dramatic. We also have found the results are relatively insensitive to the choice of n , provided that the energy loss during n time steps is still negligible.

It is important to ensure that the duplicated particles can spatially separate from each other. If the duplicated particles have the exact same position and momentum then they will never move apart and their wakes will add coherently and the stopping power will remain unchanged. Furthermore, the use of finite size particles means that particles that start off close together within a cell will move apart slowly. Several schemes can be thought of to separate the duplicated particles, such as randomly shifting the position of the new particle within the same cell, slightly shifting the momentum of the two particles in a way that conserves the total momentum (this can introduce a small divergence), or simply relying on a collisional operator, which is already used in many of the scenarios of interest, to naturally separate the two particles in phase space. We have experimented both with the momentum shifting scheme and with the use of a Monte-Carlo binary Coulomb collisions operator [55, 57, 58, 59], with overall satisfactory results.

We next show results using the OSIRIS Coulomb collisions module[59] for electron-ion collisions. We initialize a 2D plasma in a box $1\mu\text{m}$ wide and periodic in the transverse

direction, and $200\mu\text{m}$ long in the longitudinal direction. The box is filled with a $100n_c$ plasma, except for a thin ($2\mu\text{m}$) vacuum layer to the left which isolates the plasma from this wall; a laser with $I = 5 \times 10^{19} \text{ W/cm}^2$ (normalized vector potential $a_0 = 6$) and $\lambda = 1\mu\text{m}$ is incident from this direction. In order to avoid using a larger simulation box to capture the self-consistent plasma expansion due to laser heating of the plasma surface, we used infinitely heavy ions. We use a linear density ramp from 0-100 n_c in the first micron of plasma. We use cells with a size of $0.5c/\omega_p$, and either 16 or 512 particles per cell, corresponding to 64 or 2048 PPSD², respectively; the time step satisfies the Courant condition almost exactly, to reduce numerical Cerenkov. We use third order particle shapes, with 4-pass current smoothing with a compensator.

In figure 4.6 we show the particle count and forward heat-flux of a typical laser-solid simulation as a function of kinetic energy on a logarithmic scale for $N = 512$ ¹. The first thing to notice is that the majority of the energy flux is being carried by particles with $\gamma \sim 2 - 10$, demonstrating that the energy is being carried by relativistic electrons and that macro-particle stopping needs to be considered. It is also clear that the peak of the heat-flux and the peak of the particle count are separated by two orders of magnitude indicating that almost all of the forward heat flux is carried by only a few percent of the particles. This illustrates that if we decrease the particle charge and increase the particle count in only a small region of phase space, we can greatly minimize the importance of macro-particle stopping without significant computational expense. To illustrate the effectiveness of the splitting algorithm, in Figure 4.7 we show the total forward heat flux as a function of longitudinal position for the $N = 512$ case as well as a $N = 16$ case with no splitting and a $N = 16$ case with five particle splits (for final particle size equivalent to $N = 16 \times 2^5 = 512$). Figure 4.7b clearly shows significant loss of beam energy going into the plasma for the $N = 16$ case as compared to the $N = 512$ case. In fig. 4.7b the heatflux essentially vanishes by $x_1 = 100\mu\text{m}$. Taking the peak beam particle energy as $2a_0m_e c^2$ [104], we find for these parameters a stopping distance

¹ Note that these plots are of the number of particles at a given value of $\log_{10}(\gamma - 1)$, that is to say the y-value is not given by $f(\gamma - 1) = f(10^{\log_{10}(\gamma - 1)})$, but rather by $\ln(10)10^{\log_{10}(\gamma - 1)}f(\gamma - 1)$, such that the area under any section of the curve (the number of particles) is the same as $\int d(\gamma - 1)f(\gamma - 1)$ (up to a normalization factor).




FIGURE 4.7: (Color) Particle count as a function of γ and forward position for a) 2048 PPSD², b) 64 PPSD², and c) 64 PPSD² but with up to 5 particle splits for superthermal particles. d) Forward heatflux as a function of forward position for 2048 PPSD² (red curve), 64 PPSD² (magenta), 64 PPSD² w/5 splits (blue), and 64 PPSD² w/7 splits (pink). All data in arbitrary units but with equivalent normalization.

of $97.5 \mu\text{m}$, consistent with these results. For the $N = 512$ case (Fig. 4.7a), we predict a decrease of the particle Lorentz factor, γ , of 0.375 over the same distance, which we estimate to be 3–6% of beam particle energy. The larger apparent loss of heat flux moving forward in the case of $N = 512$ is due to a number of non-macro-particle effects, including the increase of laser absorption as a function of time combined with the time-of-flight of electrons, the fact that not all heat flux is being carried by highly-relativistic particles, and the effects of the wakes of at least partially coherent particle bunches. We note that collisional stopping, however, is not a significant factor in this regime. More importantly, the results using $N = 16$ but with 5 binary particle splits (Fig. 4.7c) are nearly identical to the $N = 512$ simulation. Figure 4.7d shows the comparison of the heat flux carried in the forward direction as a function of the longitudinal positions for the different cases. Again, it is possible to observe that by splitting the high-energy particles the macro-particle stopping is controlled. For the example shown, using the splitting algorithm leads to only a factor of two increase in the computational time, giving a total computational savings factor of 16 to reach the same accuracy in terms of macro-particle stopping. This illustrates the efficiency of this algorithm and its usefulness for carrying out multi-dimensional studies of large plasma volumes. We should also note that using this splitting algorithm can bring important statistical advantages when computing collisions between fast and background particles in many of the scenarios of interest. This algorithm can lead to load balancing issues. OSIRIS does have a dynamic load balancing capability and we will experiment with this as part of future work.

4.7 Summary

In this chapter we have shown that relativistic particles moving in a cold background plasma in PIC simulations are susceptible to enhanced stopping due the use of macro-particles. The stopping scales as q^2/m so that particles with large charge but with the correct charge to mass ratio will stop more rapidly. This stopping is due to the wakefield created by relativistic particles and it can be predicted using wakefield theory developed for studying plasma wakefield acceleration. We reviewed the derivation of the wakefields created by a point par-

ticle moving near the speed of light for one dimension (charge sheet), two dimensions (line of charge), and three dimensions. We used this wakefield (Green’s function) to derive the wakefields created by finite size particles (such as those used in PIC simulations), and also calculated the force on a particle from its own wakefield to get the stopping power. We also studied how the stopping depends on the cell size, particle shape, and dimensionality. The enhanced stopping is mitigated through the use of larger cells, higher order particle shapes, and current smoothing as well as with a decrease in the macro-particle charge.

We found good agreement between the theory and results against PIC simulation from a finite difference PIC code (OSIRIS) and a spectral PIC code (PARSEC). We also studied how a distribution of electrons containing a hot tail relaxes and showed that the macro-particle stopping process dominates how the tail relaxes. We also reexamined previous results in intense laser-solid interactions, such as the isolated target simulations in [65] and found that indeed this effect modified the physics. When the charge per particle was reduced by a factor of 6.25 the amount of laser energy being deposited in the core decreased by only 25% from the values quoted in Tonge et al.[65] Our expressions also predict that simulations done by other groups will also be impacted by this effect [102, 103]. Our simulations show that increasing the cell (particle) size can reduce this effect. However, for large cells the dispersion relation of the laser near the overdense plasma is also modified, so this is not necessarily a solution.

We recommend that more careful analysis be carried out on the effects of numerical dispersion of light and plasma waves, including the possibility of negative group velocity when cell size, Δ , exceeds c/ω_p . Furthermore for large cell sizes, the different effects on wakefield analysis between using finite difference operators and differential (spectral) operators to solve for the field quantities should be investigated.

We also describe a possible solution which is a particle splitting algorithm that can reduce this effect by decreasing the charge of high-energy particles as the simulation evolves. We found that this algorithm can successfully control the macro-particle stopping in PIC simulations of high-energy density scenarios in a computationally efficient way.

CHAPTER 5

Conclusion

In this dissertation we have reported on studies of three different areas of physics related to how an intense laser beam is absorbed at a sharp interface between vacuum and solid density plasma. This physics is of direct relevance to the Fast Ignition fusion concept as well as ion acceleration. These three areas are, the laser-induced acceleration of electrons at a smooth surface; the transport of finite width beams through over-dense plasma; and macro-particle stopping, a numerical effect with important relevance to the computational modeling of FI scenarios. In addition to analytical theory, all our work relied heavily on Particle-In-Cell simulation, specifically the multi-paradigm and massively parallel PIC code OSIRIS. The ongoing development of the OSIRIS codebase has allowed for an ever-increasing scope in the problems with which it can be tasked, and indeed many of the simulations described in this dissertation were computationally infeasible or even practically impossible when we first started considering them. Therefore a broader underlying lesson of this dissertation, beyond the physical insights gained, is the power of active and collaboratively developed codes, in particular OSIRIS, especially when coupled with the continuing growth of computational power, much of which is now happening in the realm of High Performance Computing (HPC).

In Chapter 2 we explored the mechanisms for converting laser energy into electron kinetic energy at a sharp vacuum to dense plasma interface in the high intensity regime. In order to focus on general mechanisms rather than specific results that arise even in a theoretical experiment with its extra complications, as well as ion acceleration, we limited our study to rather simplified geometry: either strictly one-dimensional or else two-dimensional but without any initial transverse dependence in either the matter or the light field (plane waves). In doing so, we were able to identify and explore two distinct mechanisms with starkly

different implications for the feasibility of the FI scheme. However we did learn that, although both mechanisms were present in almost all two-dimensional simulations we explored despite the initial one-dimensional nature of the profile (the exception being a step-function density gradient and a circularly polarized laser), only one mechanism was possible in the strictly 1D PIC simulations. This mechanism occurs through a simple motion in a standing wave near the (over-)critical surface, and generates accelerated electrons with momentum up to $2a_0m_e c$ (where a_0 is the normalized vector potential); we therefore call it the $2a_0$ mechanism.

In fact, it was only with an insight gained from the multi-dimensional simulations – that the high-density surface quickly heats up to relativistic temperatures – that we were able to obtain $2a_0$ acceleration in the single-dimensional code. In simulations, only when we initialize the plasma with an unrealistically high (and arbitrary) initial transverse temperature (transverse heating being [impossible in 1D](#)) was the $2a_0$ process observed in 1D. In addition, key facets of the interaction – like the absorption fraction, or the peak electron energy – were strongly dependent on this arbitrary choice of temperature, and as such the insight which could be gained from these one-dimensional simulations, especially in isolation, was strictly limited. Specifically, the 1D simulations were key to identifying how and under what conditions the $2a_0$ process occurs. However, they cannot be used to understand how a plasma can be heated and to quantify the absorption for experimental conditions. These points are important because a number of previous models based on similar one-dimensional codes struggled to accurately account for the absorption, and often arrived at erroneous explanations. However, in concert with the understanding from the 2D cases, the 1D simulations were able to demonstrate and elucidate the physical picture of the near surface acceleration in a way which we find clear and hopefully even intuitive once understood. In particular, the dual and complementary roles of the standing wave magnetic field and the transverse thermal velocity have been explained.

We found that a second mechanism, which again did not occur in 1D at all, is caused by the chaotic motion of the the electrons in the light field present in the underdense plasma, and leads to a population of electrons without a particular upper-bound on their kinetic energy. The stochastic acceleration of electrons in counter propagating lasers has been dis-

cussed before but our simulations show how an underdense plasma can form and that these electrons are produced independently of the $2a_0$ electrons. We found that the creation of this underdense plasma is related to the unevenness of the critical surface, particularly that plasma leaves the bulk in density ‘spikes’ at this surface. These spikes pierce the standing wave structure of the reflected laser and allow electrons to ‘leak out’ and enter the vacuum region. We found that this effect could be mitigated or eliminated by using circularly or elliptically polarized light.

An interesting feature of the $2a_0$ mechanism is that, when the surface remains flat, all accelerated electrons have finite transverse momentum in the direction of the laser electric field. This is caused by the need for the electrons to use their transverse momentum to be rotated out into the vacuum region by the magnetic field at the plasma surface, along with the conservation of transverse canonical momentum due to the symmetry of the flat surface. Since the momentum in the direction of the magnetic field would create no such magnetic force, the initial momentum in this direction should be irrelevant. It seems therefore that the acceleration of electrons at the surface may lead to a beam of electrons with distinct divergence in the two directions, which may be measurable in the laboratory. To explore this we ran an initially cold 3D simulation and looked at the longitudinal and transverse momentum phase spaces for the two transverse directions, which are shown in figure 5.1. As can be seen, after 50 fs (left column) the two phase spaces are indeed different. However after 100 fs (right column), the distinction has all but disappeared. It seems like this may be caused by the evolution of surface to be uneven, or due to magnetic field filaments which appear in the bulk plasma, possibly due to Weibel-like instabilities. The exact cause of this effect is an area for future research.

In Chapter 3 we investigated the transport of energetic electrons, by attempting to explain an experimental result using computer modeling. In the experiment, an electron beam created by a high intensity laser showed greatly decreased transport through a plasma layer as compared to an initially (foamed) solid. “Pure” particle-in-cell simulations were unable to fully explain the discrepancy. However similar simulations with the addition of a binary Coulomb collision model did show qualitative differences between low- and high-density




FIGURE 5.1: Momentum phase spaces for 3D simulations of $2a_0$ acceleration simulations, where the laser polarization has the electric field in the $\hat{x}_2$ direction. The left-hand column shows the results after 50 femtoseconds of laser illumination, with the top frame having p_2 on the y -axis and the bottom frame having p_3 on the y -axis (and both having p_1 on the x -axis). As was seen in two dimensional simulations, two distinct sets of accelerated electrons can be seen in the top frame, with each having either positive or negative transverse momentum. In the other transverse direction the accelerated electrons show no preference for non-zero momentum, as seen in the bottom frame. The right-hand column shows the same data after 100 femtoseconds of laser illumination, at which point the distinction between the two phase spaces is no longer visible.

transport layers, showing partial resistive collimation in the higher density plasma which was absent in the low density case (with or without collisions). Further enhancing the effect of the collisions by fixing the Coulomb logarithm at modestly high value (~ 50) lead to an even starker difference between the two cases, with the electron beam even curving back towards the axis in the dense plasma.

That collisional physics is an important aspect of the experimental outcome is not completely surprising, as both cases would be expected to have significant collisionality for the background electrons. However it was not expected that the difference between the two would be so pronounced, as the density ratio was less than a factor of seven; the initial temperature doesn't seem to explain the discrepancy either, as varying the initial temperature in the simulations doesn't seem to affect the outcome at all.

The necessity of the enhanced Coulomb logarithm is also a little unclear. Although the enhanced collisionality was certainly sufficient to collimate the beam, it may easily be that it is not totally necessary; even when the Coulomb logarithm is calculated automatically there are obvious differences in the heat flux, which in the larger geometry of the actual target may be enough to explain the experimental result. If fixing the Coulomb logarithm at a large value does have some physical validity, it probably lies in the fact that the heat capacity of the simulated materials is greatly diminished from that of the real ones, and so the enhancement somewhat corrects for the lower collision strength due to the unrealistically hot background.

In short, although the limitations of the simulation leave some quantitative questions unanswered, we find compelling albeit semi-qualitative evidence that resistive collimation is the explanation for the experimental result. Simulations which increase the density of the gold layer and increase the transverse width, or which have some method for cooling the electrons in the gold, are directions for future work.

Lastly, in Chapter 4, we looked at a numerical effect which we termed *macro-particle stopping*, which is a deceleration of relativistic particles traveling through a plasma which is caused (or, rather, greatly enhanced) by the large charge carried by simulated electrons.

In short, this effect is caused by the fact that, although the acceleration caused by a given electromagnetic field remains always proportional to the charge-to-mass ratio of the particle – and therefore unaffected by the use of macro-particles – in the case of electrostatic wakes caused by single particles the field itself is proportional to the charge alone, and so the total acceleration is given by the charge *squared* divided by the mass. For a single simulated particle representing say thousands of physical particles the effect is magnified by exactly that amount, thousands of times the real deceleration.

An important point is that, although for individual macro-particles the effect will always be unphysically large, the same does not hold for coherent bunches of particles; there the field is correctly proportional to the total charge of the bunch, and total deceleration is linear in bunch size. This is why, for instance, this effect has not been a problem in plasma wake-field accelerator simulations. Another way to state this is that the wakes of all the particles in the bunch must add coherently, and so the bunch must be smaller than the wake wavelength, and must stay so for a significant portion of the motion. For electrons accelerated in a laser-overdense matter interaction, such as is found in Fast Ignition, there is no theoretical reason to expect such a bunch structure; and indeed our simulations have found stopping proportional to the numerical weighting down to the smallest sizes we have examined. It is for this reason that this effect is commonly found in the FI literature (including our own papers), and the FI community which must be aware of and minimize its impact. It may be that the incoming electron beam will coherently bunch as it propagates in the dense core which could then lead to coherent wakes and anomalous stopping.

APPENDIX A

Conservation of Transverse Canonical Momentum

In the Lagrangian formulation of classical mechanics, the equations of motion are given by the Euler-Lagrange equation

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{q}_i} \right) - \frac{\partial \mathcal{L}}{\partial q_i} = 0 \quad (\text{A.1})$$

where the Lagrangian $\mathcal{L}$ is a function of generalized coordinates $q_i(t)$, their time derivatives $\dot{q}_i(t)$, and possibly time t explicitly.

Given any electromagnetic field which satisfies Maxwell's equations, it is possible to describe the fields in terms of a vector potential $\mathbf{A}$ and a scalar potential ϕ , where the fields themselves are given by

$$\mathbf{B} = \nabla \times \mathbf{A}, \quad \mathbf{E} = -\nabla \phi - \frac{1}{c} \frac{\partial \mathbf{A}}{\partial t} \quad (\text{A.2})$$

From these equations it is easily verified that expressing the Lagrangian in the form

$$\mathcal{L} = \frac{m\mathbf{v}^2}{2} - q\phi + q \frac{\mathbf{v} \cdot \mathbf{A}}{c} \quad (\text{A.3})$$

gives the same equations of motion as the Lorentz force for a particle of charge q and mass m

$$\dot{\mathbf{p}} = q(\mathbf{E} + \frac{\mathbf{v} \times \mathbf{B}}{c}) \quad (\text{A.4})$$

The term in the parenthesis in equation A.1 is defined as the canonical momentum. For this Lagrangian we find that it is given by

$$\frac{\partial \mathcal{L}}{\partial \dot{q}_i} = mv_i + \frac{q}{c} A_i \quad (\text{A.5})$$

¹ Relativistically this is not correct. To make it so, one can simply replace $\frac{m\mathbf{v}^2}{2}$ with $-mc^2\sqrt{1-v^2/c^2}$; or, to be expressed in a way that is manifestly Lorentz invariant, the equation should be written as $\mathcal{L}(\tau) = \frac{1}{2}mu^\mu(\tau)u_\mu(\tau) + qu^\mu(\tau)A_\mu(\tau)$, where τ is the proper time and u is the four velocity.

Looking at equation A.1 it's readily apparent that, for any coordinate for which the Lagrangian isn't explicitly dependent on, the canonical momentum associated with that coordinate is not a function of time. Therefore, for any direction in which the fields do not vary, the physical momentum of a particle in that direction is always equal to it's initial momentum, plus the value of the vector potential in that direction at its current point in space.

It is straightforward to extend A.3 and A.3 for relativistic mechanics. However, we can also arrive at the same result without using the Euler-Lagrange equation, by putting the definitions of the fields in terms of the potentials, eq. A.2, directly into the Lorentz force law, eq. A.4. Doing so, we obtain the equation

$$\dot{\mathbf{p}} = \frac{q}{c} \left(-c \nabla \phi - \frac{\partial}{\partial t} \mathbf{A} + \mathbf{v} \times (\nabla \times \mathbf{A}) \right). \quad (\text{A.6})$$

We can then use the vector identity

$$\mathbf{v} \times (\nabla \times \mathbf{A}) = \nabla (\mathbf{v} \cdot \mathbf{A}) - (\mathbf{v} \cdot \nabla) \mathbf{A} - (\mathbf{A} \cdot \nabla) \mathbf{v} - \mathbf{A} \times (\nabla \times \mathbf{v}) \quad (\text{A.7})$$

and then drop the last two terms because $\mathbf{v}$ is an independent variable and so it does not have a derivative with respect to space. Using this result we arrive at

$$\dot{\mathbf{p}} = \frac{q}{c} \left(-c \nabla \phi - \frac{\partial}{\partial t} \mathbf{A} + \nabla (\mathbf{v} \cdot \mathbf{A}) - (\mathbf{v} \cdot \nabla) \mathbf{A} \right). \quad (\text{A.8})$$

The second and fourth terms on the right hand side can be combined, using the fact that the total time derivative of the vector potential following the particle is given by $d\mathbf{A}/dt = (\partial/\partial t + \mathbf{v} \cdot \nabla)\mathbf{A}$. This term can then be moved to the left hand side to obtain

$$\frac{d}{dt} \left(\mathbf{p} + \frac{q}{c} \mathbf{A} \right) = \frac{q}{c} \nabla (\mathbf{v} \cdot \mathbf{A} - c \phi). \quad (\text{A.9})$$

Here, it can be clearly seen that if the scalar in the parentheses on the right hand side is invariant in a given direction (i.e. if it is translationally invariant), the corresponding component of the vector inside the parentheses on the left hand side, i.e. the canonical momentum, is a constant in time.

APPENDIX B

High-Energy Particle Damper in OSIRIS

B.1 Refluxing

In Chapter 2, our intent was to investigate the acceleration of electrons by a high-intensity laser at an over-dense surface. Our goal was to isolate the basic characteristics of the laser absorption in a general sense, which would hopefully be independent of the structure of the target beyond the laser interaction region. In particular, we are unable to model an entire Fast Ignition target, even in two dimensions, and so would necessarily need our results to not depend on the (arbitrary) volume of plasma we are able to simulate. Another way of stating our goal is that we would like to as closely as possible model the absorption into a plasma target of unbounded depth. Although many targets – including Fast Ignition targets[65] – will in fact not be very closely approximated by such a model, and more holistic simulations would certainly be warranted, our hope was that in separating the general from the particular, we would gain a better understanding of both.

Ideally we would simply simulate a plasma which is for our intents infinite. Unfortunately because the accelerated electrons will be traveling relativistically, we essentially need to have the back of the box causally isolated from the front for this to be the case. Increasing the duration of the simulation thus increases the computational cost quadratically, and for longer pulse lengths this quickly becomes infeasible. A finite-sized target which behaves like an infinite target is therefore necessary.

That a finite sized and unbounded target will interact with the laser beam differently is arguably intuitive. At the most basic level putting energy continuously into a real target would heat the plasma at a rate inversely proportional to the total volume. Since we have

seen that the laser absorption is highly dependent on the surface temperature, were the whole plasma to become even hotter than the surface it would certainly feed back on the absorption mechanism.

In fact however the size of the target can affect the absorption of the laser long before the material heats up appreciably, the reason for which is an effect which is called *refluxing*. This effect arises from the fact that the electrons accelerated by the laser are super-thermal, that is their energy is much much greater than the typical thermal energy in the plasma; and from the related fact that these electrons travel collisionlessly through the target.

After being accelerated, the particles move ballistically and relativistically towards the rear of the target. For a real finite target with a vacuum on the far side, the very first electrons will simply exit the target entirely and continue on in the “lab.” However this will quickly lead to a charge separation between these electrons and the ions which are left behind, which will create a strong electrostatic field at the back of the target. Within the material this field will be shorted out by the background electrons, so the high-energy particles will still travel unimpeded through the plasma; immediately upon exiting though they will experience the field and will be pulled elastically back into the target. After this reflection the fast electrons will then take the same journey in reverse, going back towards the laser interaction region along a ballistic path. When they reach the front surface again they will be able to enter the laser fields with much greater kinetic energy than the thermal energy even in the strongly heated surface layer, and they can then be re-accelerated by the laser and sent back into the plasma to repeat the cycle. This ‘bouncing’ back and forth of the energetic electrons through the target is what is meant by refluxing.

Refluxing as described so far is what would occur in a flat, wide target with a plane-wave laser. For our simulations this is the appropriate model. For a finite width laser, or say a round target, the motion of the accelerated electrons would in general be more complicated, but in many cases similar refluxing would still occur. For instance in a round target the electrons may travel not directly back towards the laser interaction region but rather along the outer surface of the target through multiple reflections.

The reflecting electrostatic field in a physical target arises because an adjacent vacuum region allows electrons to escape and leave behind ions; the exact same effect would be observed in a simulated target with the rear boundary being a vacuum of infinite or at least large enough extent. However analogous effects will also occur for most numerical boundary conditions at the rear of the target. In particular the three most common particle boundary conditions – absorbing, reflecting, and thermalizing – all lead to significant particle refluxing. How this occurs is slightly different in each case.

Reflecting boundary conditions lead to refluxing in the most straight forward manner. When particles reach the boundary their momentum remains unchanged except the component perpendicular to the wall is reversed, just like a rubber ball bouncing off a smooth wall. In this case no electrostatic field is created at the plasma edge, the particles are simply shoot back into the plasma with no mediating field. Figure B.1.b shows a phase space plot of longitudinal momentum and position for a simulation where an $a_0 = 6$ laser has been illuminating a $50\mu m$ thick target for 1 picosecond and particles are specularly reflected at the rear wall. As can be seen by comparing this plot to that in B.1.a – where the box length is equal to the simulation duration times the speed of light and so is effectively infinitely long – the reflecting boundary condition is having a significant effect on the phase space, with many more high energy particles traveling both backwards and forwards through the box in this case.

At first an absorbing boundary condition might seem like a solution to refluxing, as particles which reach the wall are deleted from the particle array and hence removed from the simulation; if there is no particle there can be no refluxing. The problem is that because PIC self-consistently solves Maxwell’s equations and the continuity equation, removing a particle doesn’t actually remove its field. The effect is more like freezing the charge at the point at which it strikes the wall, and so just like electrons escaping into vacuum an electrostatic field builds up at the back of the box from all the absorbed particles. Eventually the field is strong enough that even the hottest particles can’t reach the back of the simulation, and so are instead reflected back into the box. If mobile ions are present they can drift towards the wall as well and cancel out the charge, but in general this is far too slow to short out the

field and prevent refluxing.

Of the three particle boundary conditions at the wall, thermalizing at least in theory is the one which may be able to address refluxing. This boundary condition is similar to reflecting, but instead of re-emitting the particle with the same momentum magnitude as it struck the wall with, instead it is assigned a random energy picked from a Maxwellian distribution with a given temperature (it maintains the proper particle flux at the wall). In this case, therefore, the wall is able to remove energy without removing the particle itself. However in practice refluxing is still an issue, for the basic reason that it is not clear what the correct temperature with which to re-emit the particle is. If it is too hot, not enough energy will be removed and the situation will be similar to a reflecting boundary. Contrariwise, if it is too cold then the thermal diffusion of the cooled particles back into the plasma will be too slow to counter the flow of energetic electrons towards the wall, an electrostatic field will build up and the situation will be similar to an absorbing boundary. In figure [B.1.c](#) a phase space plot for a simulation with thermal boundary conditions at the rear of the target is shown, where the re-emission temperature is the initial temperature of the simulation, 120 eV. The results are quite similar to the reflecting boundary conditions shown in figure [B.1.b](#).

B.2 Damper

The problem with all these boundary conditions, including a vacuum ‘boundary condition’ as in a lab, is related to the fact that the plasma has to somehow generate a neutralizing return current to balance out the forward current of the accelerated electrons. It has to do this for two reasons. First, if it didn’t the net current would be the forward current alone, which greatly exceeds the Alfvén current [[63](#), [62](#), [64](#)], which is not possible (actually in this case the return current would probably be the forward current itself, turned around by its own magnetic field). Second, by accelerating electrons the laser is essentially taking particles off of the surface and shooting them into the plasma; if there were no return current the front of the target would be striped of electrons and there would be nothing left to be accelerated.

In all cases however the plasma does generate a return current. In the case of an infinite




FIGURE B.1: Phase space plots $(p_1 x_1)$ for simulations with differing boundary conditions at the rear wall of the target. In frame (a) the rear boundary is causally separated from the laser (the box length, not all shown here, is equal to c times the simulation time) and so is effectively boundary-less. Frame (b) shows the result for specular reflection at the rear of the target. For the simulation shown in frame (c) the particles are re-emitted at the rear wall with a Maxwellian distribution with a temperature of 120 eV (the initial temperature of the simulation). In frame (d) the super-thermal particles are damped to the local temperature throughout the rear $10\mu m$ of the plasma.

plasma, the charge separation between the accelerated electrons and the remaining ions is continuously shorted by the background plasma as the electrons travel deeper into the target. In the case of a finite target, this shorting happens until the beam electrons leave the plasma. Once they do exit the target however the vacuum acts as an insulator and is unable to short the field, and so the return current is instead given by the refluxing electrons. Similarly the numerical boundaries are unable to short the field or provide a return current, and so refluxing also occurs.

Our goal therefore in modeling an infinite target with a finite simulation is to find a way to remove energy from the energetic electrons but still provide a self-consistent return current. The way we do this is by stopping the relativistic electrons not at a sharp wall, but over an extended volume. In this way the electrostatic field is dispersed over enough material that it can drive the bulk plasma backwards without also reflecting the high gamma particles. In essence we have created an extended boundary condition for the particles.

The details of how we do so in OSIRIS are as follows. We denote some region near the back of the box to be the damping region. Within this volume, high energy particles are selected at random and thermalized to a given temperature in one time-step. We do this rather than slowing the particles down gradually¹ – which arguably is somewhat more physical – because it allows us to spread the stopping locations of the particles throughout the damper volume without any knowledge of the beam characteristics; it also prevents us from having to track particles across multiple nodes as they decelerate. The disadvantage of this method is that it will lead to the emission of bremsstrahlung radiation. However this radiation is both poorly resolved by the simulation grid in most cases and is effectively damped by the high density plasma.

At a given position, the damping of the particles is controlled by three parameters: a cut-off energy, a stopping distance, and a re-emission temperature. The cut-off energy specifies the way to distinguish between background particles (for which nothing happens) and hot electrons (for which damping may occur); particles above the energy are hot, those below are

¹ The damping region can also be a finite region in the middle of the plasma. This was done in [65] for integrated Fast Ignition simulations. In this case the electrons were slowed gradually by a drag force.

background. The stopping distance is used to calculate the probability that a stopping event occurs; using this number, the velocity of the particle, and the time between damping events (usually but not necessarily the simulation time-step) the code decides stochastically whether or not to stop the particle, such that for an ensemble of particles the mean distance they travel before being damped is given by the stopping distance. The re-emission temperature is used to define the Maxwellian from which the particle’s momentum is randomly chosen. For both the cut-off energy and the re-emission temperature, it is possible to either specify absolute numbers, or else to calculate the parameters based on the plasma temperature at the location where the particle is damped.

The damping region is parameterized by two longitudinal positions, denoted x_s (for ‘x start’) and x_m (for ‘x max’); x_s is at some location near the rear of the target, and x_m is located somewhere closer still to the rear wall. In front of x_s the damper is turned off, between x_s and x_m the stopping parameters are increased linearly, and from x_m to the back of the box the parameters stay at their maximum value. The reason we designed the damper this way is that, by starting with modest damping parameters and then increasing them to more stringent ones, we avoid any sharp transitions in the physics which might lead to wave reflections or other spurious events, without having to fine-tune our input parameters to the characteristics of a given beam.

Typical values for the stopping distance would start at two times the distance of x_s to the wall, and would end at one quarter of the distance of x_m to the wall. For the cut-off energy, typically it would start at 10 times the local thermal energy, and end at 2.5 times. Re-emission was in most cases at the local temperature everywhere. The whole damping volume would be $4 - 10\mu m$ deep.

These parameters were found to be sufficient, but we did not perform a detailed analysis to be sure that they are necessary. In particular, $10\mu m$ is likely a conservative estimate for how thick the damping region needs to be. However, even though this region could span multiple nodes this was usually a small computational cost, and so we did not try to optimize this parameter but rather erred on the side of caution and left it at this rather aggressive value.

Figure B.1.d shows the longitudinal momentum and position phase space plot for a simulation with particle damping in the rear $10\mu m$ of the plasma (included in the figure). The phase space results appear similar to the boundary-free case in frame (a), and much different than either reflecting or thermalizing boundary conditions. That the damper is likely thicker than it needs to be can be seen in that there are no energetic particles visible in the last $5\mu m$ of the box.

B.3 Fourth-root Temperature Calculation

Earlier we stated that we could use the temperature at a given location to define both the cut-off energy and the re-emission temperature for the damped particles. We didn't explain how that temperature is computed, though. The answer is that in most cases, we simply calculate the average energy per particle, and then use the relation $\langle E \rangle = 3/2T$ for a three dimensional Maxwellian to calculate the temperature.

However there is a problem with this method, which is that our distribution is distinctly not a Maxwellian, nor well approximated by one. Rather it is a high density, relatively low temperature background plasma, with a sparse but very high energy beam superimposed within it. The temperature we are trying to measure is that of the background, but it is often the case that most of the energy in the plasma is carried by the beam electrons, and so measuring the average energy across both species may greatly exaggerate this temperature. Therefore, we also implemented an option in the code to calculate the temperature in a way which is much less sensitive to a high-energy tail; the way we did that is to calculate, instead of the average energy per particle, the average value of the fourth root of the energy.

To interpret the result, we note that if the distribution function of the plasma is given by a normalized Maxwellian

$$f(p) = n_p \left(\frac{1}{2T\pi} \right)^{3/2} e^{-\frac{p^2}{2T}} \quad (\text{B.1})$$

then the integral of the fourth root of the (non-relativistic)² kinetic energy is readily found

² Note that the code does not assume the energies are non-relativistic but rather uses $\gamma - 1$ to calculate

to be

$$n_p \left(\frac{1}{2T\pi} \right)^{3/2} \int d^3p \left(\frac{p^2}{2} \right)^{\frac{1}{4}} e^{\frac{-p^2}{2T}} = n_p \frac{2}{\sqrt{\pi}} T^{\frac{1}{4}} \text{Gamma} \left(\frac{7}{4} \right) \quad (\text{B.2})$$

where $\text{Gamma}()$ is the standard gamma function; the temperature is then found by measuring n_p by an unweighted sum, dividing that out, and inverting the equation to solve for T . Closed-form solutions for one- and two-dimensions are also easily found.

If the distribution function is instead the sum of two Maxwellians, the same integral can be found to give

$$\int d^3p \left(\frac{p^2}{2} \right)^{\frac{1}{4}} (f_1 + f_2) = \frac{2}{\sqrt{\pi}} \text{Gamma} \left(\frac{7}{4} \right) (n_1 T_1^{\frac{1}{4}} + n_2 T_2^{\frac{1}{4}}) \quad (\text{B.3})$$

where $n_1 + n_2 = n_p$. Since the densities add linearly but the temperatures as the fourth root, a high temperature but low density tail will distort this result less than the linear sum.

In the end, as explained briefly in the text, when using this calculation we found the results comparable to the simple average of the particle energy. This makes sense, as the average energy would also be much less than for the beam alone, since the background is much denser, and so no matter what we would still be cooling the beam.

the sum. Because this method measures the temperature of the background plasma, which is not expected to be relativistic, this equation should still be valid.

APPENDIX C

Analytic Model for $2a_0m_e c$ Peak Momentum

If we assume a standing wave structure of perfect reflection at the surface (that is, no skin depth or absorption effects) and that the electron is relativistic (i.e. that its peak velocity is nearly c) it is possible to arrive at two relatively simple differential equations for the motion of the electron in laser fields. The vector potential for such a standing wave is given by

$$\mathbf{A} = \frac{2a_0m_e c^2}{e} \sin(k_0 x_1) \cos(\omega_0 t) \hat{x}_2 \quad (\text{C.1})$$

which can then be used in the equations of motion in terms of the potentials, eq. [A.9](#), to find the rate of change of the longitudinal momentum. Assuming the particle is massless (i.e. that its rest mass can be ignored), this rate of change is given by

$$\frac{dp_1}{dt} = \frac{c p_2}{\sqrt{p_1^2 + p_2^2}} 2a_0 m_e c k_0 \cos(k_0 x_1) \cos(\omega_0 t) \quad (\text{C.2})$$

where the position of the particle, x_1 , is itself a function of time, whose rate of change with respect to time is given simply by

$$\frac{dx_1}{dt} = \frac{c p_1}{\sqrt{p_1^2 + p_2^2}}. \quad (\text{C.3})$$

The conservation of transverse canonical momentum means the the transverse physical momentum, p_2 , is an explicit (not differential) function of longitudinal position and time, given by

$$p_2(x_1(t), t) = p_{20} - 2a_0 m_e c \sin(k_0 x_1) \sin(\omega_0 t). \quad (\text{C.4})$$

Substituting this into equations [C.2](#) & [C.3](#) we arrive at two coupled differential equations for the longitudinal position and momentum, which along with initial conditions p_{10} , p_{20} , and t_0 , can be solved numerically for the total change in forward momentum due to the interaction with the standing wave.

It is possible to derive an analytical estimate of the peak momentum of an electron accelerated in the standing wave if we make some simplifying assumptions. The first assumption, motivated by studying particle tracks for the most energetic electrons, is that the motion of the particle in the vacuum is circular, in particular that the longitudinal position is harmonic and so a sinusoidal function of time given by

$$x_1(t) = \frac{c}{\omega_0} \sin(\omega_0 t - \phi). \quad (\text{C.5})$$

We then assume that in order to find the total change in momentum it is only necessary to integrate electric force on the particle

$$\Delta p = \int dt eE \quad (\text{C.6})$$

This of course ignores the effect of the magnetic field, which is required to rotate first into the vacuum then back into the plasma. However because when the magnetic field is important the electric field is negligible – because the two fields are out of phase in both space and time – this is approximately correct.

There are some important points about equation C.5. First it says the peak velocity of the particle does not exceed c , which is good. It also says that the particle is initially moving normal to the plasma surface. This is somewhat inconsistent with the fact that the particle needs transverse initial momentum to enter the wave; however it is the case that the surface B field quickly rotates the particle out towards the vacuum, so it's an acceptable approximation. Next, it assumes that the particle spends one-half laser cycle in the laser fields; this somewhat makes sense as this is how long the laser electric field is pointed in one direction, although as we'll see in a moment the interaction is more complicated than that. Last, it assumes the the farthest the particle gets into the vacuum region is c/ω_0 . This means that even the most highly accelerated particles are not able to reach the area of peak electric field, which is $c\pi/2\omega_0$ from the surface; it is in fact necessary that this is true, as at this point the magnetic field is at a spatial node, meaning there is no way for the particle to be rotated back towards the plasma surface.

Now we simply need to motivate what the correct value of ϕ is. To do so we look at the

equations for the fields:

$$E_2(x, t) = \frac{2a_0m_e c \omega_0}{e} \sin(\omega_0 t) \sin(k_0 x) \quad (\text{C.7})$$

$$B_3(x, t) = \frac{2a_0m_e c \omega_0}{e} \cos(\omega_0 t) \cos(k_0 x). \quad (\text{C.8})$$

Looking at the electric field alone, taking $\phi = 0$ would give the maximum possible acceleration¹. However the interaction with the magnetic field is not consistent with the circular motion we have assumed, as the field would be symmetric in time and so would first rotate the particle one way then the other. Assuming $\phi = \pi/2$ we find the phase of the magnetic field is now correct for rotating the particle in a circular arc; but the electric force will now be symmetric and hence integrate to zero. We therefore take the phase to be half-way between these values, $\phi = \pi/4$. This choice makes sense in terms of the magnetic force in that when the particle exits the plasma, and has a lower total momentum, it sees a small and decreasing B -field, which rotates it out; and then after it is accelerated and its total momentum is high it uses a full 3/4's of the magnetic half-cycle to rotate back in. This choice is also consistent with the particle tracks.

We can now use this value of ϕ to calculate the total acceleration of the particle. The result is

$$\begin{aligned} & 2a_0m_e c \int_{\pi/4\omega_0}^{5\pi/4\omega_0} dt \omega_0 \sin(\omega_0 t) \sin(\sin(\omega_0 t - \pi/4)) \\ &= 2a_0m_e c \int_0^\pi d\tau \sin(\tau + \pi/4) \sin(\sin(\tau)) \\ &= \pi J_1(1) \sqrt{2} a_0 m_e c \end{aligned} \quad (\text{C.9})$$

where $J_1()$ is the first Bessel function of the first kind. Numerically this value is approximately $1.955a_0m_e c$.

¹ $\sim 2.76a_0m_e c$

REFERENCES

- [1] Max Tabak, James Hammer, Michael E. Glinsky, William L. Kruer, Scott C. Wilks, John Woodworth, E. Michael Campbell, and Michael D. Perr. Ignition and high gain with ultrapowerful lasers. *Physics of Plasmas*, 1(5):1626–1634, May 1994.
- [2] 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. *Lecture Notes in Computer Science*, 2331:342, 2002.
- [3] R. G. Hemker, F. S. Tsung, V. K. Decyk, W. B. Mori, S. Lee, and T. Katsouleas. Development of a parallel code for modeling plasma based accelerators. *Proceedings of the 1999 Particle Accelerator Conference*, 1999.
- [4] Fu Qiang. Radiation (solar). In James R. Holton, editor, *Encyclopedia of atmospheric sciences*, volume 5, pages 1859–1863. Academic Press, 2003.
- [5] Wayne R. Meier. Inertial fusion reactor designs. *Journal of Fusion Energy*, 7(2):183–189, 1988.
- [6] Stefano Atzeni and Jürgen Meyer ter Vehn. *The Physics of Inertial Fusion*. Oxford Science Publications, 2004.
- [7] John M. Dawson. Particle simulation of plasmas. *Rev. Mod. Phys.*, 55:403–447, Apr 1983.
- [8] C. K. Birdsall and A. B. Langdon. *Plasma Physics via Computer Simulation*. Taylor and Francis, 1st edition, 1991.
- [9] A. Henig, S. Steinke, M. Schnürer, T. Sokollik, R. Hörlein, D. Kiefer, D. Jung, J. Schreiber, B. M. Hegelich, X. Q. Yan, J. Meyer-ter Vehn, T. Tajima, P. V. Nickles, W. Sandner, and D. Habs. Radiation-pressure acceleration of ion beams driven by circularly polarized laser pulses. *Phys. Rev. Lett.*, 103:245003, Dec 2009.
- [10] 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. *Phys. Rev. Lett.*, 108:235004, Jun 2012.
- [11] J. Ongena, R. Koch, R. Wolf, and H. Zohm. Magnetic-confinement fusion. *Nat Phys*, 12(5):398–410, 05 2016.
- [12] R. Betti and O. A. Hurricane. Inertial-confinement fusion with lasers. *Nat Phys*, 12(5):435–448, 05 2016.
- [13] "Tokamak". <https://www.merriam-webster.com/dictionary/tokamak>, June 1 2017.
- [14] K. Ikeda. Progress in the iter physics basis. *Nuclear Fusion*, 47(6), 2007.

- [15] Stephen A. Slutz and Roger A. Vesey. High-gain magnetized inertial fusion. *Phys. Rev. Lett.*, 108:025003, Jan 2012.
- [16] J D Lawson. Some criteria for a power producing thermonuclear reactor. *Proceedings of the Physical Society. Section B*, 70(1):6, 1957.
- [17] R. Betti, P. Y. Chang, B. K. Spears, K. S. Anderson, J. Edwards, M. Fatenejad, J. D. Lindl, R. L. McCrory, R. Nora, and D. Shvarts. Thermonuclear ignition in inertial confinement fusion and comparison with magnetic confinement. *Physics of Plasmas*, 17(5):058102, 2010.
- [18] M. J. Edwards, J. D. Lindl, B. K. Spears, S. V. Weber, L. J. Atherton, D. L. Bleuel, D. K. Bradley, D. A. Callahan, C. J. Cerjan, D. Clark, G. W. Collins, J. E. Fair, R. J. Fortner, S. H. Glenzer, S. W. Haan, B. A. Hammel, A. V. Hamza, S. P. Hatchett, N. Izumi, B. Jacoby, O. S. Jones, J. A. Koch, B. J. Kozioziemski, O. L. Landen, R. Lerche, B. J. MacGowan, A. J. MacKinnon, E. R. Mapoles, M. M. Marinak, M. Moran, E. I. Moses, D. H. Munro, D. H. Schneider, S. M. Sepke, D. A. Shaughnessy, P. T. Springer, R. Tommasini, L. Bernstein, W. Stoeffl, R. Betti, T. R. Boehly, T. C. Sangster, V. Yu. Glebov, P. W. McKenty, S. P. Regan, D. H. Edgell, J. P. Knauer, C. Stoeckl, D. R. Harding, S. Batha, G. Grim, H. W. Herrmann, G. Kyrala, M. Wilke, D. C. Wilson, J. Frenje, R. Petrasso, K. Moreno, H. Huang, K. C. Chen, E. Giraldez, J. D. Kilkenny, M. Mauldin, N. Hein, M. Hoppe, A. Nikroo, and R. J. Leeper. The experimental plan for cryogenic layered target implosions on the national ignition facility—the inertial confinement approach to fusion. *Physics of Plasmas*, 18(5):051003, 2011.
- [19] John D. Lindl, Peter Amendt, Richard L. Berger, S. Gail Glendinning, Siegfried H. Glenzer, Steven W. Haan, Robert L. Kauffman, Otto L. Landen, and Laurence J. Suter. The physics basis for ignition using indirect-drive targets on the national ignition facility. *Physics of Plasmas*, 11(2):339–491, 2004.
- [20] J. Meyer-ter Vehn. On Energy Gain of Fusion Targets. *Nuclear Fusion*, 22(4):561–566, January 1982.
- [21] I. V. Igumenshchev, W. Seka, D. H. Edgell, D. T. Michel, D. H. Froula, V. N. Goncharov, R. S. Craxton, L. Divol, R. Epstein, R. Follett, J. H. Kelly, T. Z. Kosc, A. V. Maximov, R. L. McCrory, D. D. Meyerhofer, P. Michel, J. F. Myatt, T. C. Sangster, A. Shvydky, S. Skupsky, and C. Stoeckl. Crossed-beam energy transfer in direct-drive implosions. *Physics of Plasmas*, 19(5):056314, 2012.
- [22] R.H. Lehmberg and S.P. Obenschain. Use of induced spatial incoherence for uniform illumination of laser fusion targets. *Optics Communications*, 46(1):27 – 31, 1983.
- [23] M. Tabak, D. S. Clark, S. P. Hatchett, M. H. Key, B. F. Lasinski, R. A. Snavely, S. C. Wilks, R. P. J. Town, R. Stephens, E. M. Campbell, R. Kodama, K. Mima, K. A. Tanaka, S. Atzeni, and R. Freeman. Review of progress in Fast Ignition. *Physics of Plasmas*, 12(5):057305, 2005.

- [24] I. N. Ellis. *Wakes in Inertial Fusion Plasmas*. PhD thesis, University of California, Los Angeles, 2014.
- [25] S. Rand. Inverse Bremsstrahlung with High-Intensity Radiation Fields. *Physical Review*, 136(1B):B231–B237, October 1964.
- [26] V. P. Silin. Nonlinear High-Frequency Plasma Conductivity. *JETP*, 20(6):1510–1516, June 1965.
- [27] John F. Seely and Edward G. Harris. Heating of a Plasma by Multiphoton Inverse Bremsstrahlung. *Physical Review A*, 7(3):1064–1067, March 1973.
- [28] L. Schlessinger and J. Wright. Inverse-bremsstrahlung absorption rate in an intense laser field. *Physical Review A*, 20(5):1934–1945, November 1979.
- [29] A. Bruce Langdon. Nonlinear Inverse Bremsstrahlung and Heated-Electron Distributions. *Physical Review Letters*, 44(9):575–579, March 1980.
- [30] M. Moll, M. Schlages, Th. Bornath, and V. P. Krainov. Inverse bremsstrahlung heating beyond the first Born approximation for dense plasmas in laser fields. *New Journal of Physics*, 14(6):065010, June 2012.
- [31] C.D. Decker, W.B. Mori, and J.M. Dawson. Nonlinear collisional absorption in laser-driven plasmas. *Physics of Plasmas*, 1(12):4043–4049, December 1994.
- [32] J. P. Freidberg, R. W. Mitchell, R. L. Morse, and L. I. Rubsinski. Resonant Absorption of Laser Light by Plasma Targets. *Physical Review Letters*, 28(13):795–799, March 1972.
- [33] James Albritton and Paul Koch. Cold plasma wavebreaking: Production of energetic electrons. *Physics of Fluids*, 18(9):1136–1139, September 1975.
- [34] William Case. Parametric instability: An elementary demonstration and discussion. *American Journal of Physics*, 48(3):218–221, March 1980.
- [35] Bedros B. Afeyan and Edward A. Williams. Unified theory of stimulated raman scattering and two-plasmon decay in inhomogeneous plasmas: High frequency hybrid instability. *Physical review letters*, 75(23):4218, 1995.
- [36] D W Forslund, J M Kindel, and E L Lindman. Theory of stimulated scattering processes in laser-irradiated plasmas. *The Physics of Fluids*, 18(8):1002–1016, August 1975.
- [37] David S Montgomery. Two decades of progress in understanding and control of laser plasma instabilities in indirect drive inertial fusion. *Physics of Plasmas*, 23(5):055601, May 2016.
- [38] S. C. Wilks, W. L. Kruer, M. Tabak, and A. B. Langdon. Absorption of ultra-intense laser pulses. *Physical Review Letters*, 69:1383–86, 1992.

- [39] R. Kodama, P. A. Norreys, K. Mima, A. E. Dangor, R. G. Evans, H. Fugita, Y. Kitagawa, K. Krushelnick, T. Miyakoshi, N. Miyanaga, T. Norimatsu, S. J. Rose, T. Shozaki, K. Shigemori, A. Sunahara, M. Tampo, K. A. Tanaka, Y. Toyama, T. Yamanaka, and M. Zepf. Fast heating of ultrahigh-density plasma as a step towards laser fusion ignition. *Nature*, 412:798–802, 2001.
- [40] Paul Gibbon. *Short pulse laser interactions with matter*. World Scientific Publishing Company, 2004.
- [41] M. Borghesi, J. Fuchs, S. V. Bulanov, A. J. Mackinnon, P. K. Patel, and M. Roth. Fast ion generation by high-intensity laser irradiation of solid targets and applications. *Fusion Science and Technology*, 49(3):412–439, 2006.
- [42] Paul Gibbon and A. R. Bell. Collisionless absorption in sharp-edged plasmas. *Physical Review Letters*, 68(10):1535–1538, January 1992.
- [43] Scott C. Wilks and William L. Kruer. Absorption of Ultrashort, Ultra-Intense Laser Light by Solids and Overdense Plasmas. *IEEE Journal of Quantum Electronics*, 33(11):1954–1968, November 1997.
- [44] F. Brunel. Not-So-Resonant, Resonant Absorption. *Physical Review Letters*, 59(1):52–55, July 1987.
- [45] W. L. Kruer and Kent Estabrook. JXB Heating by very intense laser light. *Physics of Fluids*, 28(1):430–432, January 1985.
- [46] A. J. Kemp, Y. Sentoku, and M. Tabak. Hot-electron energy coupling in ultraintense laser-matter interaction. *Physical Review E*, 79(6):066406, June 2009.
- [47] J. Sanz, A. Debayle, and K. Mima. Model for ultraintense laser-plasma interaction at normal incidence. *Physical Review E*, 85(4):046411, April 2012.
- [48] A. Kemp and L. Divol. Interaction Physics of Multipicosecond Petawatt Laser Pulses with Overdense Plasma. *Physical Review Letters*, 109(19):195005, November 2012.
- [49] Roy Gerrit Hemker. Particle-in-cell modeling of plasma-based accelerators in two and three dimensions. *arXiv:1503.00276 [physics.comp-ph]*, 2015.
- [50] G. A. Bird. Direct Simulation and the Boltzman Equation. *The Physics of Fluids*, 13(11):2676–2681, October 1970.
- [51] Joseph D. Huba. *NRL Plasma Formulary*. Naval Research Laboratory, Washington, D.C., 2016.
- [52] A. G. Peeters and D. Strintzi. The Fokker-Planck equation, and its application in plasma physics. *Annalen der Physik*, 17(2-3):142–157, February 2008.
- [53] E. M. Lifshitz and L. P. Pitaevskii. *Physical Kinetics*. Pergamon Press, 1981.

- [54] Marshall N. Rosenbluth, William M. MacDonald, and David L. Judd. Fokker-Planck Equation for an Inverse-Square Force. *Physical Review*, 107(1):1–6, July 1957.
- [55] Tomonori Takizuka and Hirotada Abe. A binary collision model for plasma simulation with a particle code. *Journal of Computational Physics*, 25(3):205–219, January 1977.
- [56] K. Nanbu. Theory of cumulative small-angle collisions in plasmas. *Physical Review E*, 55(4):4642–4652, April 1997.
- [57] K. Nanbu and S. Yonemura. Weighted Particles in Coulomb Collision Simulations Based on the Theory of a Cumulative Scattering Angle. *Journal of Computational Physics*, 145:639–654, August 1998.
- [58] 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.
- [59] F. Peano, M. Marti, L. Silva, and G. Coppa. Statistical kinetic treatment of relativistic binary collisions. *Physical Review E*, 79(2):025701, February 2009.
- [60] 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, 2012.
- [61] F. Fiúza and P. Alves. Private Communication, 2016.
- [62] J. R. Davies. Alfvén limit in fast ignition. *Physical Review E*, 69(6):065402, June 2004.
- [63] Hannes Alfvén. On the Motion of Cosmic Rays in Interstellar Space. *Physical Review*, 55(5):425–429, March 1939.
- [64] J. R. Davies. The Alfvén limit revisited and its relevance to laser-plasma interactions. *Laser and Particle Beams*, pages 1–12, May 2006.
- [65] J. Tonge, J. May, W. B. Mori, F. Fiuza, S. F. Martins, R. A. Fonseca, L. O. Silva, and C. Ren. A simulation study of fast ignition with ultrahigh intensity lasers. *Physics of Plasmas*, 16(5):056311, 2009.
- [66] Nicholas A. Krall and Alvin W. Trivelpiece. *Principles of Plasma Physics*. McGraw-Hill Inc, 1973.
- [67] A. R. Bell and R. J. Kingham. Resistive Collimation of Electron Beams in Laser-Produced Plasmas. *Physical Review Letters*, 91(3):035003, July 2003.
- [68] Hong-bo Cai, Shao-ping Zhu, X T He, Si-zhong Wu, Mo Chen, Cangtao Zhou, Wei Yu, and Hideo Nagatomo. Magnetic collimation of fast electrons in specially engineered targets irradiated by ultraintense laser pulses. *Physics of Plasmas*, 18(2):023106, February 2011.

- [69] S Chawla, M-S Wei, R Mishra, K U Akli, C D Chen, H S McLean, A Morace, P K Patel, H Sawada, Y Sentoku, R B Stephens, and F N Beg. Effect of Target Material on Fast-Electron Transport and Resistive Collimation. *Physical Review Letters*, 110(2):025001, January 2013.
- [70] A P L Robinson and M Sherlock. Magnetic collimation of fast electrons produced by ultraintense laser irradiation by structuring the target composition. *Physics of Plasmas*, 14(8):083105, August 2007.
- [71] Y Sentoku, E d’Humières, L Romagnani, P Audebert, and J Fuchs. Dynamic Control over Mega-Ampere Electron Currents in Metals Using Ionization-Driven Resistive Magnetic Fields. *Physical Review Letters*, 107(13):135005, September 2011.
- [72] A A Solodov, K S Anderson, R Betti, V Gotcheva, J Myatt, J A Delettrez, S Skupsky, W Theobald, and C Stoeckl. Integrated simulations of implosion, electron transport, and heating for direct-drive fast-ignition targets. *Physics of Plasmas*, 16(5):056309, May 2009.
- [73] A R Bell and J R Davies. Magnetic field in short-puls high-intensity laser-solid experiments. *Physical Review E*, 58(2):2471–2473, August 1998.
- [74] J R Davies, A R Bell, and M Tatarakis. Magnetic focusing and trapping of high-intensity laser-generated fast electrons at the rear of solid targets . *Physical Review E*, 59(5):6032–6036, May 1999.
- [75] Laurent Gremillet, Guy Bonnaud, and François Amiranoff. Filamented transport of laser-generated relativistic electrons penetrating a solid target. *Physics of Plasmas*, 9(3):941–948, March 2002.
- [76] G. E. Kemp, A. Link, Y. Ping, D. W. Schumacher, R. R. Freeman, and P. K. Patel. Coupling of laser energy into hot-electrons in high-contrast relativistic laser-plasma interactions. *Physics of Plasmas*, 20(3):033104, 2013.
- [77] Kane S. Yee. Numerical Solution of Initial Boundary Value Problems Involving Maxwell’s Equations in Isotropic Media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, March 1966.
- [78] E. L. Lindman. “Free-Space” Boundary Conditions for the Time Dependent Wave Equation. *Journal of Computational Physics*, pages 1–13, January 1975.
- [79] John Villasenor and Oscar Buneman. Rigorous charge conservation for local electromagnetic field solvers. *Computer Physics Communications*, 69(2-3):306–316, April 1992.
- [80] E. S. Weibel. Spontaneously growing transverse waves in a plasma due to anisotropic velocity distribution. *Physical Review Letters*, 2:83, 1959.
- [81] A Kemp, Y Sentoku, and M Tabak. Hot-Electron Energy Coupling in Ultraintense Laser-Matter Interaction. *Physical Review Letters*, 101(7):075004, August 2008.

- [82] A. B. Mikhailovskii. *Theory of Plasma Instabilities Volume 1: Instabilities of a Homogeneous Plasma*. Consultants Bureau, 1974.
- [83] Y Wan, C H Pai, C J Zhang, F Li, Y P Wu, J F Hua, W Lu, Y Q Gu, L O Silva, C Joshi, and W B Mori. Physical Mechanism of the Transverse Instability in Radiation Pressure Ion Acceleration. *Physical Review Letters*, 117(23):234801, November 2016.
- [84] Luís O Silva, Michael Marti, Jonathan R Davies, Ricardo A Fonseca, Chuang Ren, Frank S Tsung, and Warren B Mori. Proton Shock Acceleration in Laser-Plasma Interactions. *Physical Review Letters*, 92(1):015002, January 2004.
- [85] T. Baeva, S. Gordienko, A. P. L. Robinson, and P. A. Norreys. Zero vector potential mechanism of attosecond absorption in strongly relativistic plasmas. *arXiv*, page 1009.0848v1 (7pp), September 2010.
- [86] T. Baeva, S. Gordienko, A. P. L. Robinson, and P. A. Norreys. The zero vector potential mechanism of attosecond absorption. *Physics of Plasmas*, 18(5):056702, 2011.
- [87] M. Vranic, T. Grismayer, J. L. Martins, R. A. Fonseca, and L. O. Silva. Particle merging algorithm for PIC codes. *Computer Physics Communications*, 191:65–73, June 2015.
- [88] D. W. Forslund, J. M. Kindel, W. B. Mori, C. Joshi, and J. M. Dawson. Two-dimensional simulations of single-frequency and beat-wave laser-plasma heating. *Phys. Rev. Lett.*, 54:558–561, Feb 1985.
- [89] J.T. Mendonça and F. Doveil. Stochasticity in plasmas with electromagnetic waves. *Journal of Plasma Physics*, 28(3):485–493, 1982.
- [90] J. T. Mendonça. Threshold for electron heating by two electromagnetic waves. *Phys. Rev. A*, 28:3592–3598, Dec 1983.
- [91] A. Sorokovikova, A.V. Arefiev, C. McGuffey, B. Qiao, A. P. L. Robinson, M. S. Wei, H. S. McLean, and F. N. Beg. Generation of superponderomotive electrons in multipicosecond interactions of kilojoule laser beams with solid-density plasmas. *Phys. Rev. Lett.*, 116:155001, Apr 2016.
- [92] Hui Xu, Zheng-Ming Sheng, Jie Zhang, and Xian-Tu He. Stochastic heating and acceleration of electrons by high-intensity lasers in inhomogeneous plasmas. *Physica Scripta*, 77(4):045502, 2008.
- [93] J. A. Stamper, K. Papadopoulos, R. N. Sudan, S. O. Dean, E. A. Mclean, and J. M. Dawson. Spontaneous Magnetic Fields in Laser-Produced Plasmas. *Physical Review Letters*, 26(17):1012–1015, April 1971.
- [94] J. A. Stamper and B. H. Ripin. Faraday-Rotation Measurements of Megagauss Magnetic Fields in Laser-Produced Plasmas. *Physical Review Letters*, 34(3):138–141, January 1975.

- [95] D. W. Forslund and J. U. Brackbill. Magnetic-Field-Induced Surface Transport on Laser-Irradiated Fields. *Physical Review Letters*, 48(23):1614–1617, June 1982.
- [96] H. Sawada, T. Yabuuchi, S. P. Regan, K. Anderson, M-S Wei, R. Betti, J. Hund, M. H. Key, A. J. Mackinnon, H. S. McLean, R. R. Paguio, P. K. Patel, K. M. Saito, R. B. Stephens, S. C. Wilks, and F. N. Beg. Diagnosing laser-driven, shock-heated foam target with Al absorption spectroscopy on OMEGA EP. *HEDP*, 8(2):180–183, June 2012.
- [97] T. Yabuuchi, H. Sawada, S. P. Regan, K. Anderson, M-S Wei, R. Betti, J. Hund, M. H. Key, A. J. Mackinnon, H. S. McLean, R. R. Paguio, P. K. Patel, K. M. Saito, R. B. Stephens, S. C. Wilks, and F. N. Beg. Temporally resolved characterization of shock-heated foam target with Al absorption spectroscopy for fast electron transport study. *Physics of Plasmas*, 19(9):092705, 2012.
- [98] H. K. Chung, M. H. Chen, W. L. Morgan, Y. Ralchenko, and R. W. Lee. FLYCHK: Generalized population kinetics and spectral model for rapid spectroscopic analysis for all elements. *High Energy Density Physics*, 1(1):3–12, December 2005.
- [99] P. B. Radha, V. N. Goncharov, T. J. B. Collins, J. A. Delettrez, Y. Elbaz, V. Yu. Glebov, R. L. Keck, D. E. Keller, J. P. Knauer, J. A. Marozas, F. J. Marshall, P. W. McKenty, D. D. Meyerhofer, S. P. Regan, T. C. Sangster, D. Shvarts, S. Skupsky, Y. Srebro, R. P. J. Town, and C. Stoeckl. Two-dimensional simulations of plastic-shell, direct-drive implosions on OMEGA. *Physics of Plasmas*, 12(3):032702, 2005.
- [100] C. Stoeckl, J. A. Delettrez, R. Epstein, G. Fiksel, D. Guy, M. Hohenberger, R. K. Jungquist, C. Mileham, P. M. Nilson, T. C. Sangster, M. J. Shoup, and W. Theobald. Soft x-ray backlighting of direct-drive implosions using a spherical crystal imager on OMEGA. *Review of Scientific Instruments*, 83(10):10E501, 2012.
- [101] P. Beiersdorfer, M. J. May, J. H. Scofield, and S. B. Hansen. Atomic physics and ionization balance of high-Z Ions: Critical ingredients for characterizing and understanding high-temperature plasmas. *HEDP*, 8(3):271–283, September 2012.
- [102] B. Chrisman, Y. Sentoku, and A. J. Kemp. Intensity scaling of hot electron energy coupling in cone-guided fast ignition. *Physics of Plasmas*, 15, 2008.
- [103] A. J. Kemp, B. I. Cohen, and L. Divol. Integrated kinetic simulation of laser-plasma interactions, fast-electron generation, and transport in fast ignition. *Physics of Plasmas*, 17, 2010.
- [104] J. May, J. Tonge, F. Fiuza, R. A. Fonseca, L. O. Silva, C. Ren, and W. B. Mori. Mechanism of generating fast electrons by an intense laser at a steep overdense interface. *Phys. Rev. E*, 84:025401, Aug 2011.
- [105] Anatoly Spitkovsky. On the structure of relativistic collisionless shocks in electron-ion plasmas. *The Astrophysical Journal Letters*, 673(1):L39, 2008.

- [106] S. F. Martins, R. A. Fonseca, L. O. Silva, and W. B. Mori. Ion dynamics and acceleration in relativistic shocks. *The Astrophysical Journal Letters*, 695(2):L189, 2009.
- [107] A. G. R. Thomas, M. Tzoufras, A. P. L. Robinson, R. J. Kingham, C. P. Ridgers, M. Sherlock, and A. R. Bell. A review of Vlasov-Fokker-Planck numerical modeling of inertial confinement fusion plasma. *Journal of Computational Physics*, 231(3):1051–1079, February 2012.
- [108] M. Tzoufras, A. Tableman, F. S. Tsung, W. B. Mori, and A. R. Bell. A multi-dimensional Vlasov-Fokker-Planck code for arbitrarily anisotropic high-energy-density plasmas. *Physics of Plasmas*, 20(5):056303, 2013.
- [109] J. Tonge, J. May, W. B. Mori, F. Fiuza, and L. O. Silva. Anomalous stopping due to macro-particles in pic simulations. In *Proceedings of the 40th Annual Anomalous Absorption Conference*, 2010.
- [110] 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. *arxiv:1308.2628*, 08 2013.
- [111] T. Katsouleas, S. Wilks, P. Chen, J. M. Dawson, and J. J. Su. Beam loading in plasma accelerators. *Particle Accelerators*, 22:81, 1987.
- [112] W. B. Mori, J. M. Dawson, C. Joshi, T. Katsouleas, J. J. Su, and S. Wilks. The plasma wakefield accelerator. In *Proc. SPIE*, volume 0873, pages 214–222. SPIE, 1988.
- [113] M. Abramowitz and I. Stegun. *Handbook of mathematical functions*. Dover, 1970.
- [114] W. Lu, C. Huang, M. M. Zhou, W. B. Mori, and T. Katsouleas. Limits of linear plasma wakefield theory for electron or positron beams. *Physics of Plasmas*, 12(6):063101 (8pp), 2005.
- [115] R. A. Fonseca, S. F. Martins, L. O. Silva, J. W. Tonge, F. S. Tsung, and W. B. Mori. One-to-one direct modeling of experiments and astrophysical scenarios: pushing the envelope on kinetic plasma simulations. *Plasma Physics and Controlled Fusion*, 50(12):124034 (9pp), 2008.
- [116] R. A. Fonseca, J. Vieira, F. Fiuza, A. Davidson, F. S. Tsung, W. B. Mori, and L. O. Silva. Exploiting multi-scale parallelism for large scale numerical modelling of laser wakefield accelerators. *Plasma Physics and Controlled Fusion*, 55(12):124011, 2013.
- [117] V. Decyk. Simulation of microscopic processes in plasma. In A. G. Sitenko, editor, *Proceedings, International Conference on Plasma Physics (Kiev, USSR)*, volume 2, page 1075. World Scientific, Singapore, 1987.
- [118] T. Grismayer, J. E. Fahlen, V. K. Decyk, and W. B. Mori. On the time-dependent resonant width for landau damping: theory and pic simulation. *Plasma Physics and Controlled Fusion*, 53(7):074011, 2011.
- [119] A. Pukhov. Three-dimensional electromagnetic relativistic particle-in-cell code VLPL (Virtual Laser Plasma Lab). *Journal of Plasma Physics*, 61(3):425–433, April 1999.

- [120] R. Lehe, A. Lifschitz, C. Thaury, V. Malka, and X. Davoine. Numerical growth of emittance in simulations of laser-wakefield acceleration. *Physical Review Special Topics - Accelerators and Beams*, 16(2):021301, February 2013.
- [121] Peicheng Yu, Xinlu Xu, Viktor K. Decyk, Weiming An, Jorge Vieira, Frank S. Tsung, Ricardo A. Fonseca, Wei Lu, Luis O. Silva, and Warren B. Mori. Modeling of laser wakefield acceleration in lorentz boosted frame using em-pic code with spectral solver. *Journal of Computational Physics*, 266:124 – 138, 2014.
- [122] Peicheng Yu, Xinlu Xu, Viktor K. Decyk, Frederico Fiuza, Jorge Vieira, Frank S. Tsung, Ricardo A. Fonseca, Wei Lu, Luis O. Silva, and Warren B. Mori. Elimination of the numerical cerenkov instability for spectral em-pic codes. *Computer Physics Communications*, 192:32 – 47, 2015.
- [123] Peicheng Yu, Xinlu Xu, Adam Tableman, Viktor K. Decyk, Frank S. Tsung, Frederico Fiuza, Asher Davidson, Jorge Vieira, Ricardo A. Fonseca, Wei Lu, Luis O. Silva, and Warren B. Mori. Mitigation of numerical cerenkov radiation and instability using a hybrid finite difference-fft maxwell solver and a local charge conserving current deposit. *Computer Physics Communications*, 197:144 – 152, 2015.
- [124] Peicheng Yu, Xinlu Xu, Asher Davidson, Adam Tableman, Thamine Dalichaouch, Fei Li, Michael D. Meyers, Weiming An, Frank S. Tsung, Viktor K. Decyk, Frederico Fiuza, Jorge Vieira, Ricardo A. Fonseca, Wei Lu, Luis O. Silva, and Warren B. Mori. Enabling lorentz boosted frame particle-in-cell simulations of laser wakefield acceleration in quasi-3d geometry. *Journal of Computational Physics*, 316:747 – 759, 2016.
- [125] Peicheng Yu. *Lorentz boosted frame simulation technique in Particle-in-cell methods*. PhD thesis, UCLA, 2016.
- [126] J. Tonge. *Particle in Cell Simulation of Instabilities in Space and Astrophysical Plasmas*. PhD thesis, UCLA, 2002.