MASTER THESIS

RELATIVISTIC FLUID SIMULATION OF PLASMA-VACUUM EXPANSION WITH LTE CLOSURE

Leon Buikstra

APPLIED PHYSICS
LASER PHYSICS AND NONLINEAR OPTICS

EXAMINATION COMMITTEE
Prof. Dr. K.J. Boller (LPNO)
Dr. F.A. van Goor (LPNO)
Dr. M.M.J. Dhallé (EMS)
Dr. M. Bussmann (Helmholtz-zentrum Dresden-Rossendorf)

HZDR

HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

UNIVERSITY OF TWENTE.

JULY 2013

Contents

1 Introduction

1.1 Motivation

High energy particle beams have many applications, both in fundamental research and medical and industrial sectors, among which is cancer treatment through radiation therapy using ions, which shows big potential for reducing the negative side effects of current radiation therapy methods. The idea of accelerating ions from solid targets, often a metal or plastic foil, using high power lasers has been around since the 1970s, however proton energies in MeV range have not been achieved until after the advent of chirp-amplified lasers in the 1990s. Even the current record of 67.5 MeV protons [10] is still not enough for the 200+ MeV ions needed for cancer treatment, and the past decades have seen a lot of research, both theoretical and experimental, into increasing the maximum ion energy and understanding the mechanisms involved in the acceleration process [12]. Several processes for accelerating ions from solid targets have been proposed, the most important being Target Normal Sheath Acceleration.

1.1.1 Target Normal Sheath Acceleration

The most common mechanism for ion acceleration in solid targets is the Target Normal Sheath Acceleration (TNSA) process [12]. A high powered laser pulse is focused onto a solid target, normally a metal or plastic foil. This causes a pulse of electrons and ions to be launched from the back of the target, in a direction mostly normal to the back surface. What happens in between is a three step process. The first is the laser impact. As the laser impinges on the target, the atoms of the target are ionized and the electrons are accelerated in the field of the laser towards the back of the target. The second stage is electron transport. The heated electrons move through the target, ionizing it throughout. After this, the target is ionized across its entire depth and no longer behaves like a neutral solid. Instead, the target has locally been turned into a plasma. In the third stage, the electrons that are ejected from the back of the target are slowed down or reversed by the field of the positive ions that are still present. This forms a sheath of negative charge on the backside of the target, and an electrical field in which the positive ions on the back surface are accelerated in the direction normal to the surface. Effectively, this is a plasma expansion in which the high temperature electrons transfer their energy to the cold ions. This expansion takes place on the nanosecond timescale. In order to improve the efficiency of TNSA, and to transport the generated particle beam with conventional accelerator structures, it is necessary to understand the entire process and find predictions for the maximum ion energy and the energy spectrum as a function of the input parameters such as the power and shape of the laser pulse and the geometrical properties of the target. Numerical models implemented as computer simulations play an important role in obtaining this understanding.

1.1.2 Limits on Particle In Cell simulation

A lot of understanding of what happens during TNSA comes from simulations using the Particle In Cell (PIC) method [7]. PIC uses a kinetic model; it tracks

the motion of groups of charged particles and their interplay with the electromagnetic field: the motion of the particles generates a current influencing the electric and magnetic fields, and the fields accelerate the charged particles by the Lorentz force. Although this approach to simulating plasma effects is effective in providing a wealth of information about short processes, divergence in the total energy puts a limit on the timescales that can feasibly be simulated [7]. As a result, PIC can normally provide information only about processes in the femto- and picosecond range. This falls two or three orders of magnitude short of the nanosecond range required to simulate a full expansion up to the point where the particle beam is detected.

1.1.3 Long timescale plasma expansion models

To bridge the gap between what can be simulated with PIC, and what is actually measured in a laboratory, plasma-vacuum expansion models can be used. In fact several plasma-vacuum expansion models have been adapted to investigate the expansion phase of TNSA. These models generally rely on many assumptions, like relativistic effects not playing a role and the entire target being in thermal equilibrium, making them less generally valid. Additionally, these models tend to only describe the expansion in one dimension, in spite of the initial laser focus size normally being of the order of micrometers and the resulting beam travelling meters before being used or detected. Most of these models are difficult to apply in two or three dimensions. In order to better understand the plasma expansion process on the timescale between the initial picosecond scale and investigate effects that affect the final ion and electron spectra, a model is needed that allows simulating the expansion in multiple dimensions at this timescale.

1.2 Thesis outline

Neither kinetic numerical modeling nor widely available analytical models currently provide sufficient information on the timescale of interest for the plasma expansion phase in TNSA. In order to fill this gap, this thesis introduces a fluid approach to modeling plasmas that is easily implemented in simulations of arbitrary dimensions and can be used to simulate long timescale effects. The model has been implemented in a simulation that can simulate plasma expansion in arbitrary dimensionality. Section 2.1 introduces the general concepts of modeling collisionless plasmas as intersecting fluids. First, the general description of collisionless plasmas by the Vlasov equation is given in section 2.1.1. Next, it is shown how equations describing the evolution of ensemble quantities like particle density, momentum density and energy density can be derived from the Vlasov equation in section 2.1.2. Next, the mechanics of TNSA are introduced along with its relation to laser intensity. Next, existing models for the plasma expansion phase are introduced, along with their assumptions and limitations. Finally, the theory section shows how the early stages of TNSA can be simulated using existing PIC codes. The methods section explains how the simulation using the LTE closure is built up in full. Section 4 discusses the verification tests performed on the simulation code. Finally, in section 5 the simulation is used to investigate a simple one dimensional-expansion in comparison to the commonly used isothermal model. The differences between two-dimensional expansion and one-dimensional expansion are shown.

2 Theory

A plasma is a globally charge-neutral collection of particles containing a significant amount of free charged particles, both positive and negative. It is globally neutral in the sense that the total charge of the entire plasma is much smaller than the number of charged particles. Plasmas do allow for local charge imbalances that result in a restoring electric field, and it is these charge imbalances and the coupling to the electromagnetic fields that make plasmas behave so differently from other states of matter. Throughout the universe, plasmas appear to be much more common than the three classical states of matter with which we are familiar on earth. All stars in particular, including our sun, are made of matter in the plasma state. It is both the plasma state's unique behaviour and its ubiquitousness, as well its association with an energy barrier above that of gas formation, the ionization energy, that leads to it being interpreted as a fourth state of matter [17].

Plasmas display what is called collective behaviour. Instead of being affected primarily by adjacent particles, the particles in a plasma affect each other at a distance through the electric and magnetic fields. For many plasma states, collisions between particles can even be ignored and these distant effects dominate completely.

The following section gives an introduction to the equations necessary to understand plasma behaviour, as well as the properties of plasmas relevant to this thesis.

2.1 Fluid description of plasmas

2.1.1 Vlasov equation

A very general description of the behaviour of collisionless plasmas is given by the Vlasov equation. This equation lies at the base of many practical models of plasmas. For a single particle species i with rest mass and charge m_i and q_i, it is equal to the collisionless Boltzmann equation:

$$\frac{\partial f_i(\vec{x},\vec{p},t)}{\partial t} + \vec{v}(\vec{p}) \cdot \frac{\partial f_i}{\partial \vec{x}} + \vec{F}_L(\vec{v}(\vec{p})) \cdot \frac{\partial f_i}{\partial \vec{p}} = 0 \tag{1}$$

Where the force $\vec{F}_L$ is the Lorentz force: $F_L = q_i(\mathbf{E} + v \times B)$ and $\vec{v}(\vec{p})$ is the velocity of a single particle as a function of its momentum: $\vec{v}(\vec{p}) = \vec{p}/(m_i\gamma) = \vec{p}c/\sqrt{m_i^2 + p^2}$. The full description of the plasma is given by the Vlasov equation for each species along with Maxwell's equations:

$$\begin{aligned} \nabla \times \mathbf{E} &= -\frac{d\mathbf{B}}{dt} & \nabla \times \mathbf{B} &= \frac{1}{c^2}\frac{d\mathbf{E}}{dt} + \mu_0 J \\ \nabla \cdot \mathbf{E} &= 0 & \nabla \cdot \mathbf{B} &= \rho \\ \rho &= \sum_i q_i \int_{\vec{p}} f_i d^3\vec{p} & J &= \sum_i q_i \int_{\vec{p}} f_i \vec{v} d^3\vec{p} \end{aligned} \tag{2}$$

2.1.2 Fluid equations

The Vlasov equation gives the evolution of a plasma in 6-dimensional phase space. As a result, the equation is too complex to be solved or even numerically

approximated directly for all but a few special cases. The Vlasov equation can, however, be used to derive simpler fluid plasma models. These models do not describe the full phase space, but instead describe densities in space by taking moments of the Vlasov equation to momentum [3]:

$$\int \vec{p}^n \left(\frac{\partial f(\vec{x}, \vec{p}, t)}{\partial t} + \vec{v} \cdot \frac{\partial f}{\partial \vec{x}} + \vec{F}(\vec{v}) \cdot \frac{\partial f}{\partial \vec{p}} \right) d^3\vec{p} = 0 \tag{3}$$

where care must be taken that $\vec{p}^n$ is an nth order tensor. This means that, for example, the zeroth moment is just the integral of the Vlasov equation over momentum:

$$\int \left(\frac{\partial f(\vec{x}, \vec{p}, t)}{\partial t} + \vec{v} \cdot \frac{\partial f}{\partial \vec{x}} + \vec{F}(\vec{v}) \cdot \frac{\partial f}{\partial \vec{p}} \right) d^3\vec{p} = 0 \tag{4}$$

The first term simply becomes the change in local particle density, while the second term becomes the divergence of the particle flux: $\vec{v}_{avg} n$. The third term turns out to be zero, making clear the zeroth moment is the continuity equation.

$$\frac{\partial n}{\partial t} + \nabla \cdot (\vec{v}_{avg} n) = 0 \tag{5}$$

Like the zeroth order moment is associated with with the density evolution, the first moment gives the change in momentum density. This provides the average velocity in the zeroth moment, but includes an unknown pressure term. The second moment gives the time derivative of the stress tensor, which provides the pressure term, but contains further unknowns. Since every new moment introduces new unknowns, this approach does not lead to a closed system of equations. To produce a closed system, the value of the remaining unknowns has to be chosen based on physical arguments. As an example of such a closure, if it's known beforehand that the velocity spread does not significantly affect the phenomenon being investigated, the system of equations can be closed by setting the stress tensor equal to zero, and the fluid equations consist of only the continuity equation and a simplified momentum transfer equation:

$$\left(\frac{\partial n \vec{p}_{avg}}{\partial t} + \frac{\partial \vec{v} n \vec{p}}{\partial \vec{x}} + \vec{F}(\vec{v}) \cdot \frac{\partial f}{\partial \vec{p}} \right) = 0 \tag{6}$$

which is called the cold plasma closure.

2.1.3 Equation of state

As mentioned in the previous section, closures require equating unknown microscopic ensemble quantities to an expression of known macroscopic quantities. This is precisely what an equation of state is, and such an equation of state can be chosen based on for example empirical considerations. For example, if the plasma components are known to locally behave as a classical ideal gas in local equilibrium, the pressure term that appears in the first and second moments can be taken to be the famous equation of state of an ideal gas: $P = nRT$, where P is the pressure, n is the local density, R the ideal gas constant and T the local temperature. By employing the relation between the internal energy and temperature $\rho_i nt = \frac{3}{2} nRT$ where ρ_{int} is the local internal energy density of the fluid, the local pressure can be determined as a function of the local internal energy: $P = \frac{2}{3} \rho_i nt$. With such an expression for the pressure, the system of equations is closed. What equation of state can be used will depend on the precise properties of the plasma investigated.

2.1.4 Waves in plasmas

The dynamics of a collisionless plasma are mostly the result of distant particle interactions. These lead to waves and wave-like effects. With the cold plasma closure and the assumption that the ions remain static, the Maxwell equations can be used to derive the general wave equation for plasmas [11]:

$$\nabla^2 E - \nabla(\nabla \cdot E) = \frac{1}{c^2}\frac{\partial^2 E}{\partial t^2} + \mu_0 \frac{n_e e^2}{m_e} E \tag{7}$$

This equation allows both longitudinal waves and transverse waves. Filling in the general expression for a longitudinal wave $\mathbf{E} = \mathbf{E_0}\vec{z}sin(kz - \omega t)$ gives the dispersion relation for density waves in a plasma:

$$\omega^2 = \omega_p^2 = \mu_0 \frac{n_e e^2 c^2}{m_e} \tag{8}$$

Since the frequency of these oscillations is always equal to the electron plasma frequency and does not depend on the length of the wave vector, the group velocity of these so called Langmuir waves is zero. As a result, plasmas can sustain local oscillations. Filling in a transverse EM wave $\mathbf{E} = \mathbf{E_0}\vec{x}sin(zr - \omega t)$ gives the dispersion relation for EM radiation:

$$\omega^2 = c^2 k^2 + \mu_0 c^2 \frac{n_e e^2}{m_e} = c^2 k^2 + \omega_p^2 \tag{9}$$

showing light with a frequency lower than the plasma frequency cannot propagate. Plasmas that are too dense for a certain frequency to propagate in are called overdense. External light of this frequency sent into an overdense plasma will be reflected as if the plasma were a metal.

2.1.5 Characteristic length and timescales

The most interesting timescale of plasma dynamics is of course determined by the plasma frequency. The most interesting length scale inside a plasma is, for purposes of this thesis, the Debye length. When a charged particle is introduced in a plasma, it will attract particles of opposite charge while repelling particles of like charge, which results in a shielding effect. The Debye length is the length scale over which the field of the foreign particle is reduced by a factor e. Since on short timescales, only the electrons react quickly enough to contribute to this shielding effect, this length depends on the plasma's electron temperature and electron density:

$$\lambda_{De} = \left(\frac{\varepsilon_0 k_B T_e}{n_{e0} e^2}\right)^{1/2} \tag{10}$$

In a plasma expansion, the characteristic speed is the ion acoustic velocity:

$$c_s = \sqrt{k_b T / m_i} \tag{11}$$

2.2 External Forces on Plasmas

In practice, one way to influence a plasma externally is by intense EM radiation, like a high powered laser.

2.2.1 Shape of a Gaussian Pulse

The most common approximation of the shape of a real laser pulse is that of a Gaussian pulse. A Gaussian pulse can be described as [**?**]

$$\mathbf{E} = \mathbf{real}\left(\mathbf{E}_{env}\psi\right) = E_0\mathbf{real}\left(\sqrt{\frac{w_0}{w(z)}}e^{-\left(\frac{r}{w(z)}\right)^2}\frac{1}{\sqrt{\pi}\tau}e^{-\xi/c^2/\tau^2}e^{-i\frac{\Omega_0}{c}\xi - i\frac{1}{2}\mathbf{arctan}\frac{z}{z_R}}\right)$$
$$\xi = \left(z - ct + \frac{r^2}{2R(z)}\right) \qquad w(z) = w_0\left(1 + (z/z_R)^2\right)$$
$$z_R = w_0^2\Omega_0/2c \qquad R(z) = z\left(1 + (z_R/z)^2\right) \tag{12}$$

where Ω_0 is the laser frequency and w_0 is the size of the laser focus. The distance z_R is the Rayleigh length, which is a measure for the speed at which the laser defocuses, $w(z)$ is the width of the beam a distance z along the propagation axis and $R(z)$ is the radius of curvature at that distance. The important thing to note is that the pulse can be described by the product of a pulse envelope and a phase. The pulse's envelope $\mathbf{E}_{env}$ consists of the first two exponents, while the third and fourth are the pulse's phase ψ. The peak intensity is normally described in terms of the parameter a_0, given by $a_0 = \frac{eE_0}{m_e c\omega_0}$ [11].

2.2.2 Ponderomotive force

Intense non-uniform EM waves, like a focused high power-laser pulse, can influence a plasma, creating local charge imbalances. When an electron in a plasma is in a weak or uniform alternating field, it will just move back and forth with the frequency of the field. On average the electron does not accelerate in any direction. However, if it moves such a large distance with each oscillation that the field is significantly weaker on one extreme of the oscillation, the returning force will be too weak to fully reverse the electron's velocity. As a result, on average over multiple field oscillations, the electron will have a net acceleration towards areas with lower field intensity. This acceleration is associated with a pseudoforce called the ponderomotive force. Under the assumption that the envelope of the field only changes slowly, the ponderomotive force is given, in first-order approximation, by [11]:

$$F_p = -\frac{e^2q^2}{4\omega^2 m}\nabla\mathbf{E}_{env}^2 \tag{13}$$

The most important feature is that the force is inversely proportional to the mass of the particle, and the acceleration inversely proportional to the square of the mass. Even the lightest ions, protons, are accelerated nearly a factor five million times slower by this force. As a result, a laser beam reflecting off a solid surface will generally only impart energy to the electrons.

2.2.3 Target Normal Sheath Acceleration

The possibility of accelerating ions to many MeVs of energy from the back of a thin foil target using a highly energetic short laser pulse has been successfully shown experimentally in 2000 [21]. In that same year, the first papers investigating the precise workings of the full process appeared [5]. The effect, known

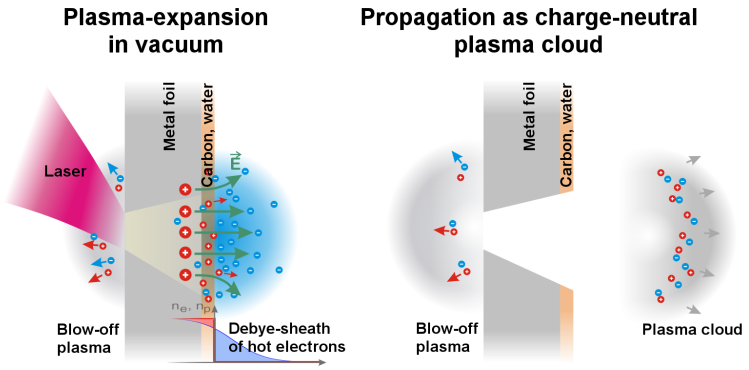

Figure 1: A cartoon showing the basic principle of TNSA acceleration. A short pulse is used to ionize a solid target, a thin foil, and heat the freed electrons. The electrons pass through the foil, ionizing it as they go. On both sides of the foil the heated electrons form a negatively charged cloud because the bulk of the electrons do not have enough energy to go further than the Debye length away from the ions. This charge cloud generates a field in which the ions are accelerated. The result is an expanding plasma, which due to the flat geometry of the foil expands primarily in the surface-normal direction. This plasma eventually forms an approximately neutral pulse of particles that propagates primarily in the target's back-surface normal direction.

as Target Normal Sheath Acceleration, consists of 3 phases. First, an intense laser pulse is focused on a thin target, often a metal foil. For low intensity, metals simply reflect laser radiation because the free electrons in the metal react quickly enough to produce a counteracting field. For high intensities however, the metal can become ionized, locally creating an expanding, hot plasma. Several mechanisms now cause the electrons to be accelerated in the laser pulse's propagation direction and into the foil. In first approximation, this can be seen as a result of a ponderomotive force, which gives both a measure for the electron energy and an explanation of why the ions are largely unaffected by the laser. Once accelerated, the electrons will pass through the target largely unimpeded. On the backside of the foil, the electrons will be ejected but will be decelerated by the field of the remaining ions to form a sheath on the back surface. The average distance the electrons can travel outside of the backside is given by the Debye length. This gives rise to the third phase, which is an expansion of the plasma on the backside of the foil into vacuum. The sheath of negative charge formed by the electrons produces a static electric field pointing in the normal direction, which accelerates the ions outwards. This in turn lowers the potential of the sheath, and as a result the electrons that reach the edge of the expanding plasma give off part of their energy, resulting in a nearly neutral plasma pushing a front of ions. After the expansion phase, the ions and accompanying electrons form a charge normal particle beam. The entire process is shown schematically in figure 1. [10]

2.2.4 Ponderomotive scaling

In the case of a thin foil, the electrons can contribute to the field accelerating the ions until they have passed their kinetic energy to the ions. In this case it can be expected that the electrons transfer a significant amount of their initial energy to the ions, and the ion energy scales with the electron energy. In first approximation this will correspond to the ponderomotive scaling. For low intensities ($a_0 << 1$) this is approximately given by the so-called ponderomotive scaling [10]:

$$E_{k,avg} = m_e c^2 \left(\sqrt{1 + \frac{a_0^2}{2}} - 1 \right) \tag{14}$$

For very high intensities ($a_0 >>$) this becomes [10]

$$E_{k,avg} = m_e c^2 \frac{\pi a_0}{ln16 + 2lna_0} \tag{15}$$

Either way the energy transferred to the electrons per affected electron monotonously increases with a_0; the highest achievable temperature depends only on the intensity, not the total energy of the laser.

2.3 Plasma expansion into vacuum

After the electrons are accelerated at the front of the target, the actual ion acceleration happens during the acceleration phase. Several models have been developed to investigate the shape of the ion spectrum. This section introduces several commonly used plasma expansion models.

2.3.1 Isothermal Mora model

The isothermal Mora model of plasma expansion [15] is a one-dimensional model based on treating the electrons as an inertialess fluid in thermal equilibrium with the ion distribution, which for determining the electron distribution is assumed to be static. This is a fluid closure where the electron momentum is set to zero and the electron pressure is set according to the ideal gas law. The electron fluid being permanently in equilibrium corresponds to the electrons having no inertia and infinite heat transfer, and their temperature being constant requires an energy source. From these conditions, the electric field is determined, which is then used as the driving force in the ion motion, which are treated using a cold fluid closure. The treatment of both fluids is classical. Under these assumptions, the electron distribution is given by the non-linear Poisson equation:

$$\varepsilon_0 \frac{\partial^2 \mathbf{\Phi}}{\partial x^2} = e(n_e - Z n_i) \tag{16}$$

where $\mathbf{\Phi}$ is the electrostatic potential, n_e is the local electron density, n_i the ion density and Z the charge number of the ions. The cold fluid closure for the ions is given by:

$$\begin{aligned} \frac{\partial v_i}{\partial t} + v_i \frac{\partial v_i}{\partial x} &= -\frac{Ze}{m_i} \frac{\partial \Phi}{\partial x} \\ \frac{\partial n_i}{\partial t} + \frac{\partial n_i v_i}{\partial x} &= 0 \end{aligned} \tag{17}$$

with v_i the local ion velocity, n_i the ion density and m_i the ion mass. By replacing the electron distribution by a quasineutrality assumption

$$n_e = Zn_i \tag{18}$$

this system has an analytical self-similar solution that predicts an infinite maximum ion energy. More recent versions of the model have modified the model by introducing an ion front at finite distance from the starting front and allowing charge separation, which requires solving the non-linear Poisson equation around the front, using the boundary conditions that the electric potential has to disappear within the foil and the electric field has to disappear far outside the foil [13]. In order to more closely resemble the conditions found in TNSA expansion, this model has frequently been modified with multiple electron species, where one species represents the cold background electrons that were present in the target initially and the hot electrons represent the electrons accelerated at the front of the foil. (for an example, see [24])

2.3.2 Adiabatic Mora model

This model is equivalent to the isothermal Mora model with two major changes. First, the initial ion density no longer occupies the half-space $x < 0$ but now occupies a symmetrical range around $x = 0$: $n_i(t = 0) = rect(x/L)$. Second, the electron temperature is now allowed to vary in time in precisely such a way that the total energy gained by the ions and the field is compensated by a loss of the electrons, keeping the total energy in the model constant:

$$\frac{\partial E_{k,e}}{\partial t} = -\frac{\partial E_{k,i}}{\partial t} - \frac{\partial E_{field}}{\partial t} \tag{19}$$

with $E_{k,e}$ the total kinetic energy of the electrons, $E_{k,i}$ the same for the ions and E_{field} the total energy stored in the field. Since the foil is assumed to be symmetrical, the boundary condition at $x = 0$ is based on this; the electric field and potential now has to be zero at the inner boundary.

2.3.3 Semi-static Schreiber model

The Schreiber model is a non-fluid model that connects the laser intensity with the maximum achievable ion energy by assuming the potential formed by the hot electrons behind the foil is static. By assuming a Maxwellian distribution, an analytical expression for the potential generated by the electrons is found. The motion of the ions in this potential, assuming they start out stationary, is then integrated over time for the duration of the pulse. The model thus gives an estimate for the highest achievable ion energy for a specific laser intensity a_0 [19].

2.3.4 Kinetic Mora model

Described in 2008, this model, like the earlier fluid expansion models from the same group, separates the electron and ion motion timescales [14]. Unlike in the isothermal and adiabatic models, no assumption is made about the global temperature or energy distribution evolution of the electrons. Instead, the electron distribution is described as a function of time and energy, assuming an

initial Maxwellian distribution. This distribution is then discretized by describing it as a number of electrons represented by their total (kinetic and potential) energy. The change in this distribution is then calculated by calculating the change in energy for each electron over the course of a round-trip through the slowly varying potential. To make sure electrons return after every round trip, the electric potential is assumed infinite far away from the foil. The electron density in space is reconstructed from these trajectories, which allows calculating a new potential; this process is finally iterated until it converges for each time step. From the electron trajectories, the local temperature can also be calculated. Very similar results can be found by assuming beforehand that the electron energy distribution is step-like [9], which is the distribution the kinetic model moves towards.

2.3.5 Assumptions for each model

The isothermal Mora model, not having any energy conservation built-in, assumes there is an energy source keeping the electrons globally at a constant temperature. The adiabatic model does not have this requirement. Both models assume that local variations in the velocity spread of the electrons do not influence the final spectra much. Similarly, barring the static Schreiber model, the models are valid only if two and three dimensional effects do not play a significant role. The kinetic Mora model requires the assumption that the electric potential goes to infinity far away from the foil, a condition that can obviously only be satisfied in one dimension.

2.3.6 Applications of each model

In the isothermal Mora model, the assumption of an energy source can be satisfied by a continuing impacting laser pulse. The general assumption in this model is that the heating starts when the heated electrons reach the backside of the target, and stops when the last electrons have returned from the backside, for a total acceleration time of the pulse duration. This means care must be taken when applying the model to thin foils, in which the electrons can contribute to the acceleration of the ions multiple times. In addition, the model says nothing about the development of the ion spectrum after the laser pulse has stopped. The adiabatic model does not have this limitation, but obviously violates causality by globally adjusting the electron temperature, invalidating the model for thick foils and late times due to the distances involved being too large. Neither model can be applied in two or three dimensions due to the temperature being independent on spatial coordinates, while in a real post-TNSA expansion the temperature will have a radial, if not more complicated dependence. The kinetic model mentioned is in principle not limited in time but cannot easily be expanded to higher dimensions. The static Schreiber model is already three dimensional, but only gives an estimate for the maximum ion energy, without any indication of the shape of the energy spectrum.

2.4 Kinetic plasma modeling: Particle In Cell model

One of the downsides of the analytical and semianalytical models referenced in the previous section is their strong assumptions on the specific situation

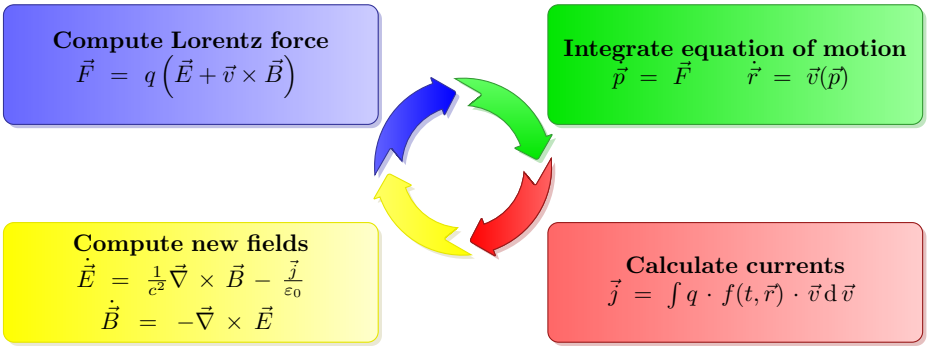

Figure 2: Time stepping in a practical numerical implementation of the PIC model.

investigated. For less specific situations a more general model is needed. Although the Vlasov equation can in principle be solved numerically directly, its 6-dimensionality in general makes this unfeasible. For short timescales, the Particle-in-Cell (PIC) model has nonetheless made it possible to efficiently simulate many plasma effects.

The premise of PIC is to approximate the phase space described in the Vlasov equation by so-called macroparticles, which have a single momentum and finite volume in space. These macroparticles represent a point symmetric distribution of real charged particles. Their time evolution is solved by computing the equation of motion of a single particle at the center of the macro-particle. The fields are calculated on a fixed grid, and currents or charge densities in the Maxwell equations are calculated from the motion or position of the macroparticles.

Because PIC simulates the full phase-space and is trivially applied in any dimensionality, it can be applied in more situations than any of the analytical models mentioned in the previous section. When implemented numerically, the PIC model puts strict requirements on the space and time resolution. On the space-side, the resolution must be high enough to resolve the Debye length as well as the smallest radiation wavelength (for example the laser wavelength), while the time step must be small enough to resolve the plasma frequency. In addition, many numerical schemes used in simulations limit the time step along with the smallest spatial resolution.

2.4.1 TNSA with PIC

The PIC model can be applied to many different situations, for example to simulate the laser impact on a foil in TNSA. It can also be used to produce reliable results of the first stages of the expansion of the formed plasma into vacuum. Figures 3 and 4 show the electron density and proton spectrum from a single two-dimensional simulation of the TNSA process performed with the mature plasma PIC code PICLS [20]. Being a single simulation, both the laser absorption, electron transport and plasma expansion are handled in the same manner. In the case shown in the figure, the laser impacted the foil under a 45-degree angle, resulting in an asymmetrical electron density distribution and an asymmetrical ion spectrum, an effect that cannot be predicted using the

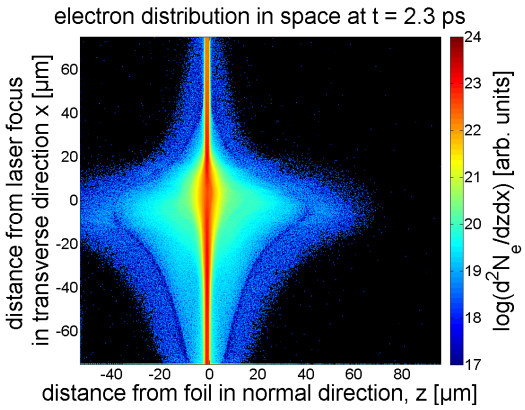

Figure 3: Electron density 2.3ps after 45 degree impact of an intense ($a_0 = 21.6$) short laser pulse (pulse duration 20 fs) on a 1 μm thin titanium foil, simulated using the particle in cell code PICLS. Due to the angle of the laser, the electron distribution is highly asymmetrical.

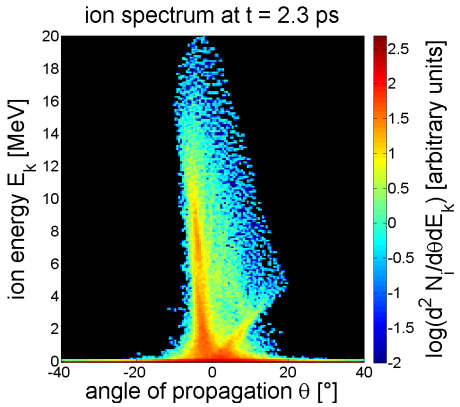

Figure 4: Spectrum of ion energy 2.3ps after laser impact under 45 degree angle. The particles whose energy spectrum is displayed in this figure are the protons that exist on the outside of the titanium foil as a result of water vapour contamination. As a result of the laser angle, the ions are ejected under an angle, but two peaks are visible. The development of this spectrum over longer timescales can be approximated using one dimensional models, but this will not provide information about the full shape of the spectrum.

one-dimensional models or static models in the previous section, although the deflection angle can be found using analytical models [26].

2.4.2 Limitation of simulations based on the PIC model

Most implementations of the PIC algorithm use explicit schemes, meaning the values for each next time step are calculated directly from the values of the current time step. Explicit schemes are generally not energy-conserving, and explicit implementations of PIC are no exception. As a result, these simulations suffer from numerical heating [6] as a result of the finite occupation of phase space. Due to the required small grid and time step size, this puts a fairly strict limit on the longest time that can be simulated using PIC algorithms. As a general matter simulating for more than a few million time steps using an explicit PIC algorithm is not possible [7]. The images shown above show that in the case of the laser coming in under an angle, a one-dimensional model will likely not be able to fully describe the resulting particle beam. This necessitates the development of new models that are not limited in time or dimensionality when the long-term effects in the plasma expansion are of interest.

3 Methods

3.1 Numerical solution of fluid plasma equations

3.1.1 Finite volume schemes

The system of equations governing fluids is of the form

$$\frac{\partial \mathbf{u}}{\partial t} - \nabla \cdot \mathbf{f}(\mathbf{u}) = \mathbf{s}(\mathbf{u}, ...) \tag{20}$$

In this, $\mathbf{u}$ is the collection of local ensemble quantities (for example density, momentum density and energy density), $\mathbf{f}$ is a function called the flux, which depends on the ensemble quantities only, and $\mathbf{s}$ is the source term which is allowed to depend on both $\mathbf{u}$ and external quantities (for example the electromagnetic field). Without the source term, the sum of the ensemble quantities over all space is conserved. The flux function represents the rate at which the quantities are transported. To conserve the quantities in a numerical simulation, most schemes for this kind of differential equation work by using the flux function to calculate a flux across cell boundaries and adjusting the densities based on the in- and outflow over the duration of a time step. If we define a volume V bounded by the surface A and call the average of the ensemble quantities in that volume at a time $t = n\Delta t$, U^n, the average at a later time $t = t + \Delta t$ is

$$\mathbf{U}^{n+1} = \mathbf{U}^n + \frac{1}{V}\int_t^{t+\Delta_t}\int_A \mathbf{f}(\mathbf{u}) \cdot dA dt = \mathbf{U}^n + \frac{1}{V}\Delta_t A \mathbf{f}_{avg} \tag{21}$$

where $\mathbf{f}_{avg}$ is the flux into the bounding surface averaged over the surface and over time. For a cubic volume with ribs Δ_x long, aligned with the axes of a cartesian coordinate system, this becomes

$$\mathbf{U}^{n+1} = \mathbf{U}^n + \frac{\Delta_t}{\Delta_x}(\vec{x} \cdot (\mathbf{f}_{x-} - \mathbf{f}_{x+}) + \vec{y} \cdot (\mathbf{f}_{y-} - \mathbf{f}_{y+}) + \vec{z} \cdot (\mathbf{f}_{z-} - \mathbf{f}_{z+})) \tag{22}$$

If the average fluxes are the exact average fluxes, this scheme is exact as well. To calculate the exact fluxes, however, the exact value of u has to be known at each point of each surface, and in a simulation, only the averages in each cell are known. To approximate the cell averages on the next time step, an approximation to the fluxes is needed, based on the known cell averages only. The big advantage of building a numerical scheme with this approach is that, since everything that leaves a cell is added to another cell, the conservative property of the original differential equation is maintained. [image of cube with fluxes] The next question is how to approximate the flux in a way that the simulation is stable and provides a way to converge to the correct solution. Several of these fluxes are introduced in the following sections and applied on the following simple, scalar, one-dimensional fluid scheme, representing a density distribution moving to the right at a constant velocity of 1/3 cells per time step:

$$\frac{\partial n}{\partial t} - \frac{1}{3}\frac{\partial n}{\partial x} = 0 \tag{23}$$

. Put into the form of equation 20, this means $\mathbf{u} = n$, $\mathbf{f}(\mathbf{u}) = \frac{1}{5}n$ and $\mathbf{s} = 0$. The examples use $\Delta x = 1$ and $\Delta t = 1$.

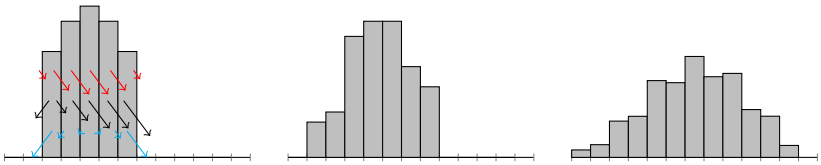

Figure 5: Lax-Friedrichs flux applied to the scalar test case. The image on the left shows the initial conditions, with the fluxes drawn in. The diffusive term is shown in cyan and the averaged flux term is shown in red. Note how the central difference term alone results in a flux out of a cell with zero density, which would result in a negative density in that cell, and an oscillation trailing behind the wave. The central image shows the state of the system after one time step, and the image on the right shows the state after 6 time steps, showing how the diffusive term smears out the distribution over time.

3.1.2 Lax-Friedrichs flux

An early, easily implemented flux that was shown to converge in one dimension for linear systems is the Lax-Friedrichs flux. Between a cell with average $\mathbf{U}_i^n$ and one with $\mathbf{U}_{i+1}^n$ the flux $\mathbf{f}_{LF}$ is approximated as []

$$\mathbf{f}_{LF}(U_i^n, U_{i+1}^n) = \frac{1}{2}(\mathbf{f}(\mathbf{U}^n i+1) + \mathbf{f}(\mathbf{U}^n i)) + \frac{1}{2}\frac{\Delta x}{\Delta t}(\mathbf{U}^n i+1 - \mathbf{U}^n i) \tag{24}$$

Effectively, the flux between the two cells is approximated by averaging over both cells. The second term represents a numerical diffusion that stabilizes the scheme. The relative size of this stabilizing term becomes bigger for smaller time steps. Figure 5 shows why oscillations will form using only the averaged central flux, and how the diffusive term dampens these oscillations.

3.1.3 Lax-Wendroff flux

A flux without added diffusion is the Lax-Wendroff flux [16]. The two-step variant, indicated here as $\mathbf{f}_{LW}(U_i^n, U_{i+1}^n)$, is:

$$\begin{aligned} \mathbf{U}_{i+1/2}^{n+1/2}(U_i^n, U_{i+1}^n) &= \frac{1}{2}(\mathbf{U}_i^n + \mathbf{U}_{i+1}^n) + \frac{1}{2}\frac{\Delta_t}{\Delta_x}(\mathbf{f}(U_i^n) - \mathbf{f}(U_{i+1}^n)) \\ \mathbf{f}_{LW}(U_i^n, U_{i+1}^n) &= \mathbf{f}(\mathbf{U}_{i+1/2}^{n+1/2}(U_i^n, U_{i+1}^n)) \end{aligned} \tag{25}$$

First, the conserved quantities are calculated at a temporary mid-point, halfway between the two cells and halfway to the next time step in time: $\mathbf{U}_{i+1/2}^{n+1/2}(U_i^n, U_{i+1}^n)$. This is done by taking the linear average of the values of the two cells as a starting point, and then using a central difference to linearly extrapolate the value at the midway point for half a time step. The flux over the entire time step is then taken to be flux at this midpoint flux. Although this produces more accurate results for linear systems of equations, it also produces oscillations that can cause instability (non-convergence or even impossibility to calculate) for non-linear systems of equations like relativistic plasmas. For example, in the simple scalar example system, it causes negative values to form

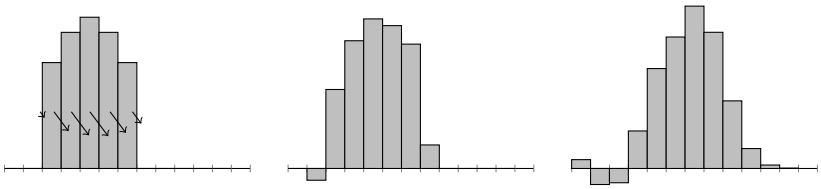

Figure 6: Lax-Wendroff flux applied to the scalar test case. The image on the left shows the initial conditions, with the fluxes drawn in. Like all higher order fluxes, oscillations form oscillations in propagation that do not show up in the analytical solution. Compared to the Lax-Friedrichs flux, the original shape of the distribution is better preserved.

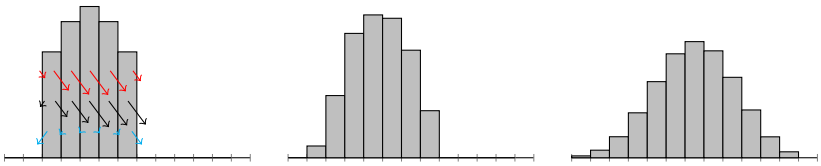

Figure 7: FORCE flux applied to the scalar test case. As for the Lax-Friedrichs case, the fluxes for the initial state are drawn in, showing the (reduced) diffusive term from the Lax-Friedrichs flux and the rest of the flux. Compared to the Lax-Friedrichs flux, the distribution is less smeared out.

behind the travelling wave. Unlike in the case of a naive central difference flux, these oscillations do not grow without bound, but negative values are still inadmissible for quantities like density.

3.1.4 FORCE flux

A more recently developed centered flux scheme is the FORCE flux (First ORder CEntered flux). Although the motivation and derivation follows an entirely different path, the flux equates to the linear average of the Lax-Friedrichs and Lax-Wendroff flux [4]:

$$\mathbf{f}_{FORCE}(U_i^n, \mathbf{U}_{i+1}^n) = \frac{1}{2}(\mathbf{f}_{LF}(\mathbf{U}_i^n, \mathbf{U}_{i+1}^n) + \mathbf{f}_{LW}(\mathbf{U}_i^n, \mathbf{U}_{i+1}^n)) \tag{26}$$

This approximation to the average flux does not form oscillations like the Lax-Friedrichs flux, but has reduced numerical diffusion compared to the Lax-Friedrich flux while also producing smoother results.

3.1.5 Flux Limited Central flux

To further reduce numerical diffusion and improve the accuracy of the solution, it is possible to construct a linear combination of a first order flux (like FORCE) and a higher order flux like the Lax-Wendroff flux in such a way that no new maxima or minima occur for the conserved quantities. One such scheme is the

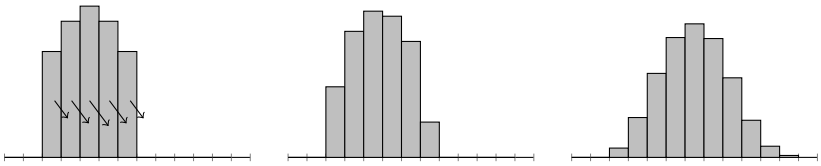

Figure 8: The upwind flux for the simple linear, scalar test case. The diffusion is even smaller than the FORCE flux, a result that could be achieved by using information about the direction of propagation of the wave.

FLIC flux (Flux LImited Centered flux) [23]. The general shape of the flux is

$$\begin{aligned}\mathbf{f}_{FLIC}(\mathbf{U}^n_{i-1},\mathbf{U}^n_i,&\mathbf{U}^n_{i+1},\mathbf{U}^n_{i+2}) = \mathbf{f}_{1o}(\mathbf{U}^n_i,\mathbf{U}^n_{i+1}) \\ &+ L(r(\mathbf{U}^n_{i-1},\mathbf{U}^n_i,\mathbf{U}^n_{i+1},\mathbf{U}^n_{i+2}))(\mathbf{f}_{2o}(\mathbf{U}^n_i,\mathbf{U}^n_{i+1}) - \mathbf{f}_{1o}(\mathbf{U}^n_i,\mathbf{U}^n_{i+1}))\end{aligned} \tag{27}$$

where $r(\mathbf{U}^n_{i-1},\mathbf{U}^n_i,\mathbf{U}^n_{i+1},\mathbf{U}^n_{i+2})$ is a function that gives a measure of the local "smoothness" of the cell averages, and $L(r)$ is a so called limiter function that determines how aggressively the second order scheme is applied. There is no general procedure to determine this smoothness, although Toro [23] suggests the following:

$$r(\mathbf{U}^n_{i-1},\mathbf{U}^n_i,\mathbf{U}^n_{i+1},\mathbf{U}^n_{i+2}) = min\left(\frac{\varepsilon^n_i - \varepsilon^n_{i-1}}{\varepsilon^n_{i+1} - \varepsilon^n_i}, \frac{\varepsilon^n_{i+1} - \varepsilon^n_i}{\varepsilon^n_{i+2} - \varepsilon^n_{i+1}}\right) \tag{28}$$

3.1.6 Non-central schemes

One feature of the centralized schemes mentioned above is that they have the same form regardless of the conservation equation being solved. The advantage of this approach is the methods are relatively easy to implement and computationally inexpensive. There is also a large amount of methods that make use of the information in the differential equation to be solved, to achieve lower amounts of numerical diffusion. For example, in the example system, the knowledge that information is travelling solely to the right allows it to be solved using an upwind scheme:

$$\mathbf{f}_{upwind}(U^n_i,\mathbf{U}^n_{i+1}) = \mathbf{f}(\mathbf{U}^n_i) \tag{29}$$

Although an upwind scheme is easily implemented for this simple equation, it becomes much more complicated for fluid equations, where for example the momentum flux resulting from pressure does not travel into a single direction. In addition to a simple upwind scheme, more complicated schemes can be devised that result in even less numerical diffusion without producing spurious oscillations. These schemes tend to be complicated and have a different shape for different sets of conservation equations.

3.2 Closures and equations of state

The two closures used for the fluid simulations in this thesis are the cold ion closure and the Local Thermal Equilibrium (LTE) closure for the electrons. The LTE closure is derived from relativistic continuum mechanics by transforming the stress-energy tensor.

3.2.1 Energy-stress tensor

In the fluid rest frame, the energy-stress tensor is given by

$$T_{\mu'\nu'} = \begin{pmatrix} \varepsilon' & & & \\ & P & & \\ & & P & \\ & & & P \end{pmatrix} \tag{30}$$

where P is the local pressure and ε' is the local energy density in this rest frame. The component $T_{\mu\nu}$ represents the flux of the μth component of the four-velocity over a surface with constant νth coordinate. For example, the component T_{00} is the energy-flux over the surface with constant time: the energy density. It is intuitive that with a source term S_ν acting on the fluid, the following must hold:

$$\partial^\mu T_{\mu\nu} = S_\nu \tag{31}$$

Applying the Lorentz transform over both indices, the above identity, written with the time and space components separated, becomes the following pair of equations:

$$\frac{\partial}{\partial t}\left(\frac{\varepsilon' + P}{1-v^2} - P\right) - \nabla \cdot \left(\vec{v}\frac{\varepsilon' + P}{1-v^2}\right) = 0 \tag{32}$$

$$\frac{\partial}{\partial t}\left(\frac{\varepsilon' + P}{1-v^2} v_i\right) - \nabla \cdot \left(\vec{v}\frac{\varepsilon' + P}{1-v^2} v_i\right) + \nabla_i P = F \tag{33}$$

In addition to these equations, we need the equation for the number density $n'/\sqrt{1-v^2}$, which is derived by taking the zeroth moment of the Vlasov equation directly:

$$\frac{\partial}{\partial t}\left(\frac{n'}{\sqrt{1-v^2}}\right) - \nabla \cdot \left(\vec{v}\frac{n'}{\sqrt{1-v^2}}\right) = 0 \tag{34}$$

Finally, we introduce the Lorentz factor $\gamma = \sqrt{\frac{1}{1-v^2}}$ and the lab-frame quantities,

$$\begin{aligned} \varepsilon &= \frac{\varepsilon' + P}{1-v^2} - P = \gamma^2(\varepsilon' + P) - P \\ \vec{p} &= \frac{\varepsilon' + P}{1-v^2}\vec{v} = (\varepsilon + \vec{p})\vec{v} \\ n &= \frac{n'}{\sqrt{1-v^2}} = \gamma n' \end{aligned} \tag{35}$$

Where $\vec{v} = \frac{\vec{p}}{\varepsilon + P}$, resulting in the following equations of motion:

$$\begin{aligned} \frac{\partial \varepsilon}{\partial t} - \nabla \cdot \vec{p} &= \vec{v} \cdot F \\ \frac{\partial p_i}{\partial t} - \nabla \cdot (\vec{v} p_i) + \nabla_i P &= F \\ \frac{\partial n}{\partial t} - \nabla \cdot (\vec{v} n) &= 0 \end{aligned} \tag{36}$$

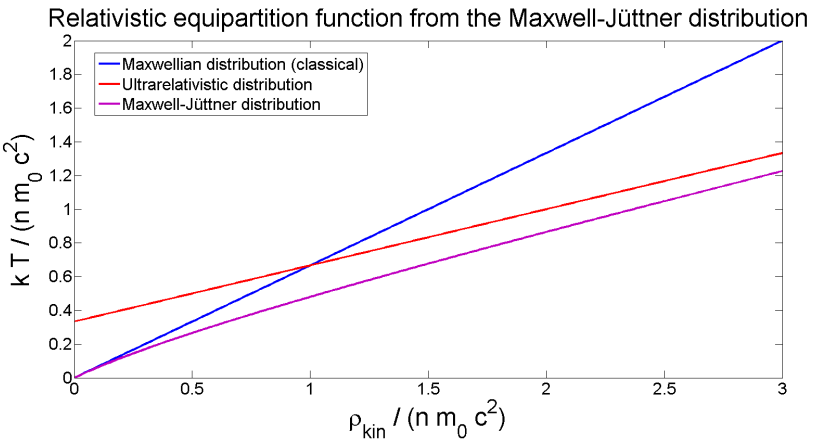

Figure 9: Relation between particle kinetic energy and pressure as predicted by the Maxwell Jüttner distribution, compared to the classical case and ultra-relativistic approximation. The temperature shows a significant deviation from either classical and ultrarelativistic approximation, making it necessary to use the full expression for the pressure found from the distribution.

3.2.2 Pressure

In order for the equations of motion found above to be a full closure, an expression for the pressure is needed. This was already implicitly assumed in the previous section, but the pressure used in this closure is isothermal. Anisotropic pressure can be included in the model by including different pressure components in the comoving frame, but this obviously makes it necessary to have three expressions for the pressure, complicating the system. Similarly, considerations have to be made about the degree with which to include relativity. For example in the non-relativistic limit, the ideal gas law gives, for single-species gases without internal degrees of freedom, a relation between the internal energy of the gas and the pressure like $\rho_{int} = \frac{2}{3}nkT = \frac{2}{3}P$. In the ultrarelativistic limit, this becomes $\rho = \frac{1}{3}nkT = \frac{1}{3}P$. Because the temperature region of interest in this thesis is neither classical nor ultrarelativistic, the pressure term needs to work in the intermediate region. The Maxwell-Jüttner distribution is the relativistic generalization of the Maxwellian distribution from which the classical equipartition for ideal gasses is derived. The equipartition theorem derived from this distribution gives the internal energy of the fluid by the implicit relation [2]

$$\frac{\rho_{int}}{n'kT} = \zeta\frac{K_3(\zeta)}{K_2\zeta} - 1 - \zeta \tag{37}$$

with $\zeta = mc^2/kT$, and ρ_{int} is the internal energy in the co-moving frame: $\rho_{int} = \varepsilon' - n'mc^2$. The function K_α is the modified Bessel function of the second kind. What this function looks like when inverted (giving the pressure as a function of co-moving internal energy) is shown in figure 9.

3.2.3 Assumptions of the LTE model

The primary assumption of the LTE closure is of course that there is a local thermal equilibrium in the initial situation. Additionally, no force gradients big enough to cause wave breaking may occur. Finally, the assumption that the plasma can be treated as a continuum has to be fulfilled. This condition is normally phrased in terms of the plasma parameter Λ, which is the number of particles in a sphere with a radius of the Debye length. The requirement is $\Lambda >> 1$. Since the Debye length depends on temperature, this condition is more easily fulfilled at high temperatures.

3.2.4 Full LTE closure

To recap, for the LTE closure, the following equations of motion are used for each fluid:

$$\begin{aligned} \frac{\partial \varepsilon}{\partial t} - \nabla \cdot \vec{p} &= \vec{v} \cdot F \\ \frac{\partial p_i}{\partial t} - \nabla \cdot (\vec{v} p_i) + \nabla_i P &= F \\ \frac{\partial n}{\partial t} - \nabla \cdot (\vec{v} n) &= 0 \end{aligned} \tag{38}$$

with

$$\begin{aligned} \gamma &= \sqrt{\frac{1}{1 - v^2}} \\ \varepsilon &= \gamma^2(\varepsilon' + P) - P \\ \vec{p} &= (\varepsilon + P)\vec{v} \\ n &= \gamma n' \end{aligned} \tag{39}$$

The defining property of the LTE closure is that the system of equations is closed by the implicit relation:

$$\frac{\varepsilon' + P}{n' m c^2} = \frac{K_3(n' m c^2 / P)}{K_2 n' m c^2 / P} \tag{40}$$

which connects the local energy density in the co-moving frame and the pressure.

3.3 Maxwell Solver

3.3.1 Yee scheme

The Yee scheme, originally published by Kane Yee in 1966 [25] is an oft used finite difference scheme to numerically solve the Maxwell equations, starting from certain initial conditions. At the base of the scheme lies approximating both the curls in the time evolution of the fields and the time derivatives themselves with central difference operators:

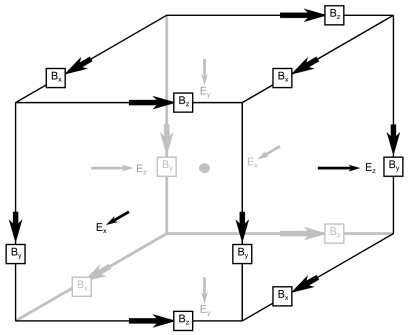

Figure 10: A Yee cell: positions on which the different field components are defined are shown.

$$\begin{aligned} \frac{\partial A}{\partial x} &\Rightarrow \frac{A(t, \vec{r} + \vec{x}\frac{1}{2}\Delta x) - A(t, \vec{r} - \vec{x}\frac{1}{2}\Delta x)}{\Delta x} \\ \frac{\partial A}{\partial t} &\Rightarrow \frac{A(t + \frac{1}{2}\Delta t, \vec{r}) - A(t - \frac{1}{2}\Delta t, \vec{r})}{\Delta t} \end{aligned} \tag{41}$$

This leads to the following time step functions:

$$\begin{aligned} \mathbf{E}_x(t + \Delta t, \vec{r}) &= \mathbf{E_x}(t, \vec{r}) + \\ &\Delta t c^2 \left(\frac{B_z(t + \frac{1}{2}\Delta t, \vec{r} + \vec{y}\frac{1}{2}\Delta_y) - B_z(t + \frac{1}{2}\Delta t, \vec{r} - \vec{y}\frac{1}{2}\Delta_y)}{\Delta y} \right) - \mu_0 J_x \\ \mathbf{B}_x(t + \Delta t, \vec{r}) &= \mathbf{B_x}(t, \vec{r}) + \\ &\Delta t \left(\frac{E_z(t + \frac{1}{2}\Delta t, \vec{r} + \vec{y}\frac{1}{2}\Delta_y) - E_z(t + \frac{1}{2}\Delta t, \vec{r} - \vec{y}\frac{1}{2}\Delta_y)}{\Delta y} \right) \end{aligned} \tag{42}$$

With the updates for E_y, E_z, B_y and B_z given by permuting the indices.

3.3.2 Staggered grid

In order to calculate the central time difference appearing in the time step function of the E-field, the value of the B-field needs to be known half a time step later. Similarly, for a time step of the B-field, the E-field is only needed half a time step after the current known state of the B-field. This allows the fields to be shifted half a time step with respect to each other to make the scheme symmetrical in time. Similarly, the central differences in space require the different components of each field to be defined at different points in space. The unit cell is shown in figure 10. This distribution of field components allows each of the central differences in equation 41 to be calculated.

3.3.3 Dispersion Relation

In order to find the dispersion relation, a complex valued test wave can be entered into the scheme. Through this method, the following relation between a wave's wave number $\vec{k}$ and its frequency ω is found [22]:

$$\frac{1}{c\Delta t}sin^2\frac{1}{2}\omega_n\Delta t = \frac{1}{\Delta x}sin^2\frac{1}{2}k_x\Delta x + \frac{1}{\Delta y}sin^2\frac{1}{2}k_y\Delta y + \frac{1}{\Delta z}sin^2\frac{1}{2}k_z\Delta z \quad (43)$$

In the case where $\Delta x = \Delta y = \Delta z = \Delta$, this is:

$$\omega = \frac{2}{\Delta t}sin^{-1}\sqrt{S^2(sin^2(\frac{1}{2}k_x\Delta) + sin^2(\frac{1}{2}k_y\Delta) + sin^2(\frac{1}{2}k_z\Delta))} \quad (44)$$

Note that this is approximately equal to the analytical dispersion relation in the case that $k\Delta << 1$ and $\omega\Delta t << 1$.

3.3.4 CourantFriedrichsLewy condition

In order for the Yee scheme to produce stable solutions, the time step size has to fulfill the CourantFriedrichsLewy (CFL) condition [22]:

$$c\Delta t \leq 1\sqrt{\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2}} \quad (45)$$

which can be derived from the fact that the no wave is allowed to grow. This means the frequency of each wave, as found in the previous subsection, needs to be real or have a negative imaginary part. In practice, for $\Delta x = \Delta y = \Delta z$, the condition simplifies to

$$c\Delta t \leq 1\sqrt{3}\Delta x \quad (46)$$

in three dimensions or

$$c\Delta t \leq 1\sqrt{2}\Delta x \quad (47)$$

in two dimensions, putting a strict limit on the largest time step that can be taken.

3.3.5 Convolutional Perfectly Matched Layer

To terminate the fields at the simulation boundaries, a so called convolutional perfectly matched layer (CPML) has been implemented. This absorbs only outgoing and incoming waves, while leaving waves travelling perpendicular to the border undamped. The concept of a perfectly matched layer is to write the Maxwell equations in the frequency domain in such a way that the wave function produces exponentially dampened waves in the absorbing layer, where the dampening is independent of frequency. When this system of equations is transformed back to the time domain, a convolution over time appears that would be prohibitively expensive to calculate if not for a trick developed by Roden and Gedney [18] that allows the convolutional terms $\psi_{E_{i,j}}$ and $\psi_{B_{i,j}}$ to be updated explicitly and inexpensively:

$$
\begin{aligned}
\psi_{E_{x,y}}(t+\frac{1}{2}\Delta t,\vec{r}) =& b_y(y)\psi_{E_{x,y}}(t-\frac{1}{2}\Delta t,\vec{r}) \\
&+\frac{c_y(y)}{\Delta y}\left(B_z(t-\frac{1}{2}\Delta t,\vec{r}+\vec{y}\Delta y)-B_z(t-\frac{1}{2}\Delta t,\vec{r}-\vec{y}\Delta y)\right) \\
\psi_{E_{x,z}}(t+\frac{1}{2}\Delta t,\vec{r}) =& b_z(z)\psi_{E_{x,z}}(t-\frac{1}{2}\Delta t,\vec{r}) \\
&+\frac{c_z(z)}{\Delta z}\left(B_y(t-\frac{1}{2}\Delta t,\vec{r}+\vec{z}\Delta z)-B_y(t-\frac{1}{2}\Delta t,\vec{r}-\vec{z}\Delta z)\right) \\
\psi_{B_{x,y}}(t+\frac{1}{2}\Delta t,\vec{r}) =& b_y(y)\psi_{E_{x,y}}(t-\frac{1}{2}\Delta t,\vec{r}) \\
&+\frac{c_y(y)}{\Delta y}\left(E_z(t-\frac{1}{2}\Delta t,\vec{r}+\vec{y}\Delta y)-E_z(t-\frac{1}{2}\Delta t,\vec{r}-\vec{y}\Delta y)\right) \\
\psi_{B_{x,z}}(t+\frac{1}{2}\Delta t,\vec{r}) =& b_z(z)\psi_{E_{x,z}}(t-\frac{1}{2}\Delta t,\vec{r}) \\
&+\frac{c_z(z)}{\Delta z}\left(E_y(t-\frac{1}{2}\Delta t,\vec{r}+\vec{z}\Delta z)-E_y(t-\frac{1}{2}\Delta t,\vec{r}-\vec{z}\Delta z)\right)
\end{aligned}
\tag{48}
$$

Updates of the other components are done similarly. Every time step, these convolutional terms are added to the fields to dampen outgoing waves [22]:

$$
\begin{aligned}
\mathbf{E}_x(t+\Delta t,\vec{r}) =& \mathbf{E}_{\mathbf{x}}(t,\vec{r})+\Delta t c^2\left(\frac{B_z(t+\frac{1}{2}\Delta t,\vec{r}+\vec{y}\frac{1}{2}\Delta_y)-B_z(t+\frac{1}{2}\Delta t,\vec{r}-\vec{y}\frac{1}{2}\Delta_y)}{\Delta y}\right) \\
&+\Delta t c^2\left(\psi_{E_{x,y}}(t+\frac{1}{2}\Delta t,\vec{r})-\psi_{E_{x,z}}(t+\frac{1}{2}\Delta t,\vec{r})\right)-\mu_0 J_x \\
\mathbf{B}_x(t+\Delta t,\vec{r}) =& \mathbf{B}_{\mathbf{x}}(t,\vec{r})+\Delta t\left(\frac{E_z(t+\frac{1}{2}\Delta t,\vec{r}+\vec{y}\frac{1}{2}\Delta_y)-E_z(t+\frac{1}{2}\Delta t,\vec{r}-\vec{y}\frac{1}{2}\Delta_y)}{\Delta y}\right) \\
&+\Delta t\left(\psi_{B_{x,y}}(t+\frac{1}{2}\Delta t,\vec{r})-\psi_{B_{x,z}}(t+\frac{1}{2}\Delta t,\vec{r})\right)
\end{aligned}
\tag{49}
$$

The coefficients a_i and b_i are zero away from the edges in the ith dimension. At the edge these values are ramped up (see [22]). Behind the absorbing layer, the edge of the simulation area is treated like a perfectly conducting mirror where the fields are kept at zero.

3.3.6 Simulation cell, current deposition, Lorentz force

The field components are defined as in 10. The center of a cell with indices i,j,k in the simulation is located at the coordinates $i\Delta x\vec{x}+j\Delta y\vec{y}+k\Delta z\vec{z}$. The x component of the electric field with these indices is shifted half a cell width down in the x direction, so $E_{x,i,j,k}$ stores, at the start of a simulation time step l, the value $E_x(l\Delta t,(i-\frac{1}{2})\Delta x\vec{x}+j\Delta y\vec{y}+k\Delta z\vec{z})$. Similarly the y component is shifted down along the y axis and the z component is shifted down in the z direction. For the magnetic field, the x component is shifted down half a cell width in both the y and z direction, the y component in x and z and the z

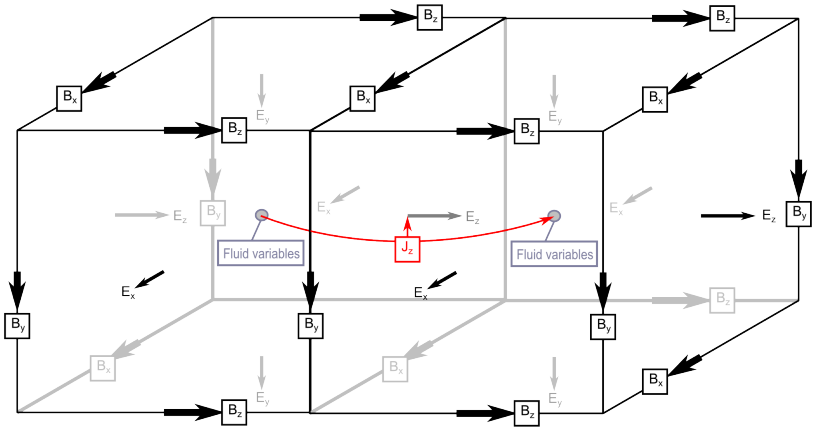

Figure 11: Two simulation cells with the flux in the z direction shown. This flux is used as the current term in the update of E_z between the two cells. When the force on the fluid is calculated, the linear interpolation of all field components surrounding the fluid variables is used.

component in x and y. The fluid components are all defined at the center of each cell. This means that the current term in the field update for $E_{z,i,j,k}$ can be calculated from the particle number flux between cells $i, j, k-1$ and i, j, k as shown in figure 11. When applying the field to a fluid cell, it has to be interpolated since the field is not known in the center of the cell. Currently, a linear interpolation is used, meaning fluid cell i, j, k feels an electric field $E = \frac{1}{2}\left(\vec{x}(E_{x,i-1,j,k} + E_{x,i,j,k}) + \vec{y}(E_{y,i,j-1,k} + E_{y,i,j,k}) + \vec{z}(E_{z,i,j,k-1} + E_{z,i,j,k})\right)$. Similarly the magnetic field is interpolated from 4 points for each component.

3.4 Parallelization

The accuracy of the approximation given by finite volume schemes depends on the grid resolution used. This means that for a certain physical size of the simulation area, in order to get a certain accuracy, a minimal amount of gridpoints in each dimension has to be calculated each time step. Due to the CFL condition, the minimal amount of steps needed for a certain simulation time increases linearly with the distance between two gridpoints. As a result, the amount of operations required for one simulation in two and three dimensions becomes too large to be processed by any single processor available. Modern supercomputers provide more processor power by having many thousands of processor cores. In order to actually use this processing capacity, calculations have to be divided into parts that can be calculated independently and therefore simultaneously.

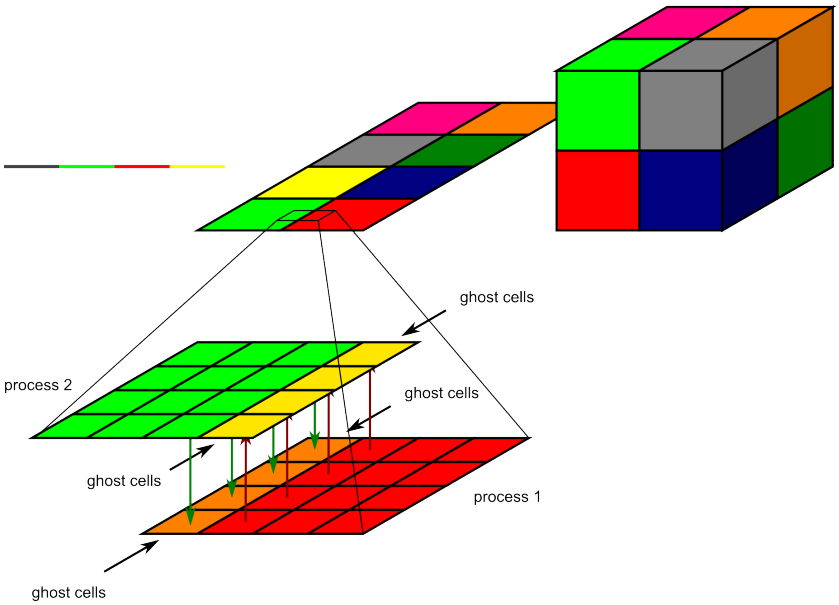

Figure 12: Domain decomposition of a one-dimensional, two-dimensional and three-dimensional simulation area. Communication of ghost cells for the two-dimensional area is shown for the case that each domain has only one layer of ghost cells.

3.4.1 Petsc

The C library Petsc [1] has been used to automate the process of dividing calculations among different processors for independent processing and coordinating the necessary communication between the processes. Although Petsc offers a large amount of mathematical functionality, only the domain decomposition features were used.

3.4.2 Domain decomposition

The way different parts of the simulation are divided over different processes is by spatially dividing the simulation grid into different domains and calculating the grid values on the next time step on a different processor for each domain. Each processor has access to the grid values it needs to calculate the values in its own domain of the simulation volume.

3.4.3 Domain geometry

The simulation area is divided into domains in all dimensions of the simulation. In one dimension, this means each processor is responsible for a consecutive series of points, in two dimensions the points are on a rectangle and in three dimensions the domain is a rectangular box. The three cases are shown in figure 12.

3.4.4 Domain size

The domains handled by each processor are sized to be as equal in size as possible.

3.4.5 Ghost cells

Because the grid values at the next time step for points at the edge of each domain depend on values existing outside the domain they cannot be calculated. To be able to calculate all new values for the next time step, the domains have a slight overlap. The cells at the edges of each domain, for which no new values are calculated, are called ghost cells. The amount of ghost cells required depends on how many values around a point are needed to calculate the values on the next time step. If a flux limited scheme is used, each domain has a two cell thick layer of ghost cells at each border. Otherwise, only one a cell layer is used. The process of copying cells from one process to the ghost cells of another process is also shown in figure 12.

3.4.6 Boundary conditions

To terminate the simulation boundaries, three sets of boundary conditions have been implemented.

3.4.7 Open boundary

For the isothermal boundaries, the fluid cells on the border of the simulation are not processed and kept at their initial values. The flux is computed as normal between the boundary cell and the internal cell, freely allowing outflux on empty border cells and influx on cells with an initial density.

3.4.8 Adiabatic/Closed boundary

For closed boundaries, the pressure and density at the border is set equal to the closest internal cell, while the momentum is mirrored over the edge. The result is that momentum can be transferred over the border, but energy and rest mass cannot. For adiabatic boundaries, no CPML layer is added.

3.4.9 Periodic boundary

For testing purposes, periodic boundaries were implemented. All grid values on the border of the simulation, including those of the fields, are set to match the values at the far end of the simulation area. In the periodic case, no CPML layers are present.

3.5 Full simulation

To recap, the equations that need to be solved are the following:

For each fluid species, the fluid equations have to be solved.

$$\begin{aligned} \frac{\partial \varepsilon}{\partial t} - \nabla \cdot \vec{p} &= \vec{v} \cdot F \\ \frac{\partial p_i}{\partial t} - \nabla \cdot (\vec{v} p_i) + \nabla_i P &= F \\ \frac{\partial n}{\partial t} - \nabla \cdot (\vec{v} n) &= 0 \end{aligned} \tag{50}$$

$$\frac{\varepsilon' + P}{n' m c^2} = \frac{K_3(n' m c^2 / P)}{K_2 n' m c^2 / P} \tag{51}$$

This is done by approximating the flux between fluid cells using the FORCE flux from section 3.1.4.

Alongside the fluid equations, the Maxwell equations have to be solved:

$$\begin{aligned} \nabla \times \mathbf{E} &= -\frac{d\mathbf{B}}{dt} & \nabla \times \mathbf{B} &= \frac{1}{c^2}\frac{d\mathbf{E}}{dt} + \mu_0 J \\ \nabla \cdot \mathbf{E} &= 0 & \nabla \cdot \mathbf{B} &= \rho \\ \rho &= \sum_i q_i \int_{\vec{p}} f_i d^3\vec{p} & J &= \sum_i q_i \int_{\vec{p}} f_i \vec{v} d^3\vec{p} \end{aligned} \tag{52}$$

This is done using the Yee scheme using CPML boundaries where possible.

Combining all of the components of the simulation results in the following time stepping procedure:

- Update CPML convolutional terms for B field using equation 48 and E-field at $t = n\Delta t$.
- Update B-field from $t = (n - \frac{1}{2})\Delta t$ to $t = (n + \frac{1}{2})\Delta t$ using equation 42 and the E-field from $t = n\Delta t$.
- Apply Lorentz force at half strength to the fluids using $\frac{\partial \rho}{\partial t} = F \cdot v$ to adjust the energy density
- Update each fluid species $t = n\Delta t$ to $t = (n+1)\Delta t$ in order, applying the flux directly as currents on the E-field. Different dimensions are currently treated in the same order every time step.
- Update CPML convolutional terms for E field using equation 48 and B-field from $t = (n + \frac{1}{2})\Delta t$.
- Complete the update of the E-field from $t = n\Delta t$ to $t = (n+1)\Delta t$ using equation 49 and the B-field from $t = (n + \frac{1}{2})\Delta t$.
- Apply the fields at half strength to the fluids as before

Between each cycle, the current simulation time is updated and optionally the state of all fluids and fields is saved. The entire scheme is implemented in C++, using Petsc functions for memory management and inter-process communication.

4 Evaluation and tests

4.1 Tests of the fluid scheme

4.1.1 Resolution and convergence using the FORCE flux

Although the FORCE-flux scheme should produce converging results, it is formally only valid for linear systems of equations. It is therefore worth verifying that the numerical implementation does indeed converge. In addition, a convergence test is a good test to show what kind of resolution is required to sufficiently reduce the effect of the numerical diffusion needed to guarantees the stability of the scheme. To this end, the one-dimensional expansion with LTE closure is repeated using a wide range of resolutions. If the resolution is sufficiently high, increasing it further should not significantly affect the density profile or proton spectrum at a specific time after the start of the simulation. Comparisons between the proton density profile and energy spectrum in a one-dimensional expansion at different grid resolutions are shown in 13.

As can be seen from the density profile, the density converges towards a single solution. As could be expected from overly large numerical dispersion, a low resolution primarily affects the area where the solution is supposed to be sharply defined, while the effect is much smaller in the area where the fluid is already expanding. One notable aspect of the spectra is the lower maximum energy at lower resolution. This is also a result of numerical diffusion. The highest proton energy occurs at a peak in the middle of the expansion profile, and the diffusion broadens and lowers this peak. Based on these tests, it seems that with the current flux scheme, for the electron temperature (1 MeV), a resolution of at least $2\mu m$ per cell is needed. At this resolution, there is still a cutoff in the maximum ion energy purely from numerical effects, but for lower energies the spectrum closely follows the curve found at higher resolutions.

4.1.2 Flux limited scheme

As higher order schemes suffer less from numerical diffusion compared to linear first order schemes, using such a scheme should improve the resolution of the simulation using the same grid size. At the time of this writing, the only higher order scheme implemented that does not suffer from strong oscillatory effects is a flux-limited total variation diminishing (TVD) scheme as described in section 3.1.5. The question is whether or not such a scheme actually improves the resolution of the simulation at the same grid size. To this end, a simple one-dimensional simulation of the expansion of an infinitely thick foil is performed using the flux limited scheme for both ions and electrons. The output of this simulation is then compared to the same initial conditions simulated with the diffusive first-order flux at different resolutions. Attention is paid to the density profile and shape of the field. Figure 14 visualizes the output of the simulation for the case of an expanding infinitely thick foil in one dimension. As shown in the figure, the scheme prevents new local minima and maxima from appearing, but does not fully prevent the numerical oscillatory effect associated with linear higher order schemes on slopes. Since ion density slopes are associated with the electric field in a plasma, this creates very large fields which can drive unphysical effects. For this reason, the flux limited scheme was not used in the final simulations.

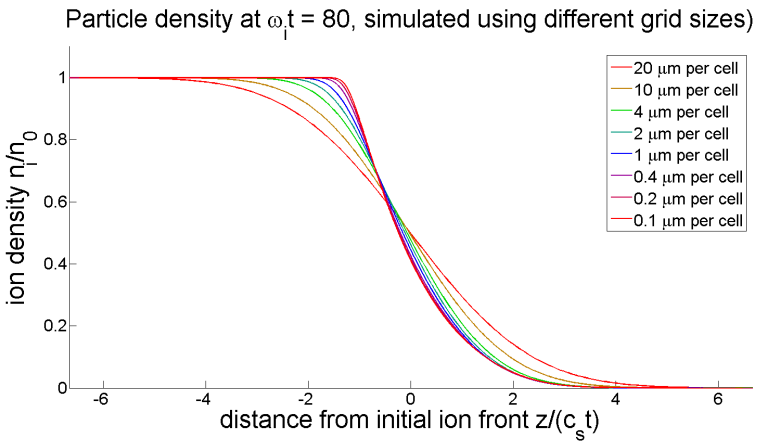

(a)

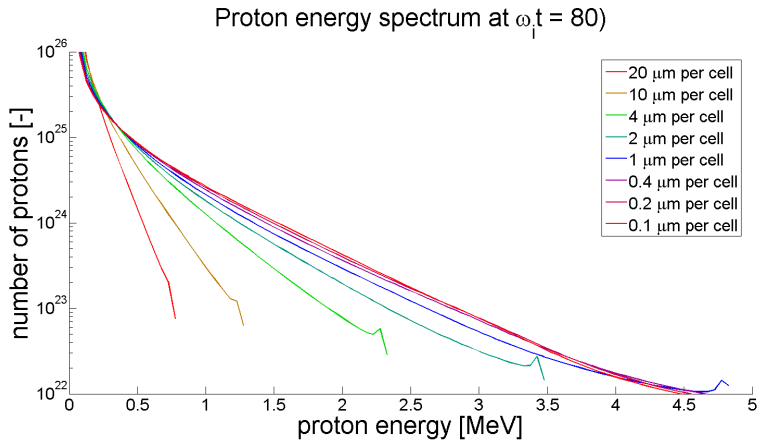

(b)

Figure 13: Comparison of the density profile and energy spectrum of protons in a one-dimensional plasma-vacuum expansion, as approximated in simulation using different grid resolutions. At high resolutions, both the evolution of the density profile and the ion energy spectrum converge to a single solution.

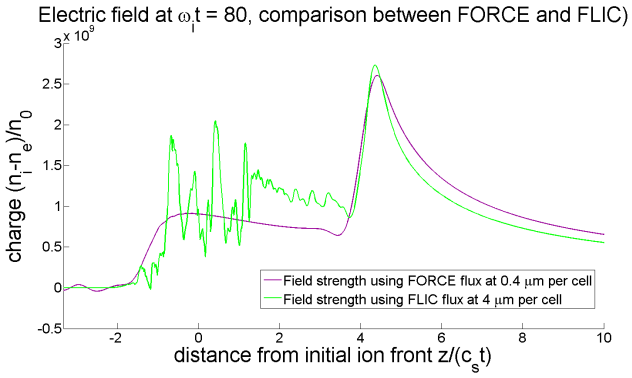

(a)

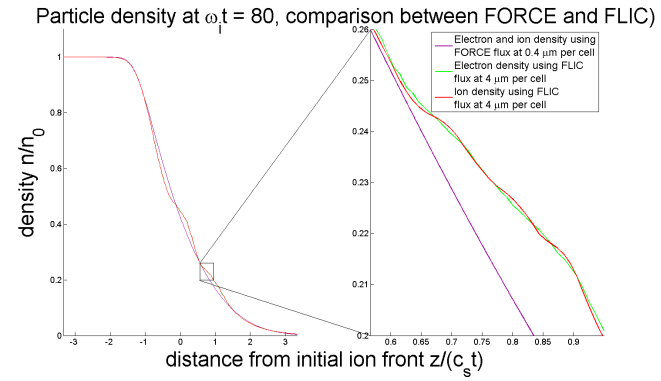

(b)

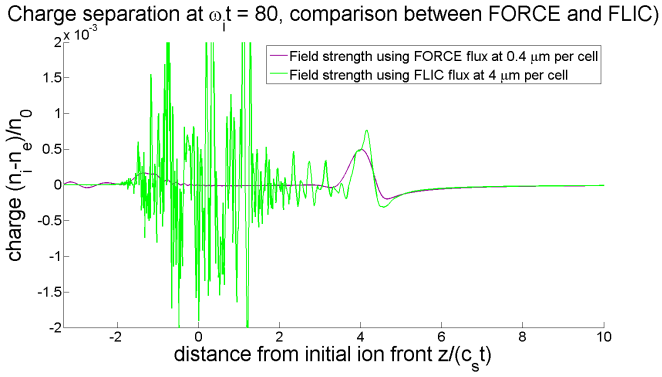

(c)

Figure 14: A comparison between expansion of an infinitely thick foil using flux limited and 1st order schemes. Although flux limiting prevents new local minima and maxima from forming, on a sloped profile this approach still allows damped oscillations to form. Although this does not result in any true numerical instability, it does result in spurious ridges on the ion density profile. Because when the electrons are close to equilibrium, the field strength is related to the steepness of the ion density distribution, this results in a noisy field which is hard to analyze; in addition it is difficult to estimate what the effect is on the final energy distribution of the ions. This seems to make flux limited schemes impractical for plasma simulations.

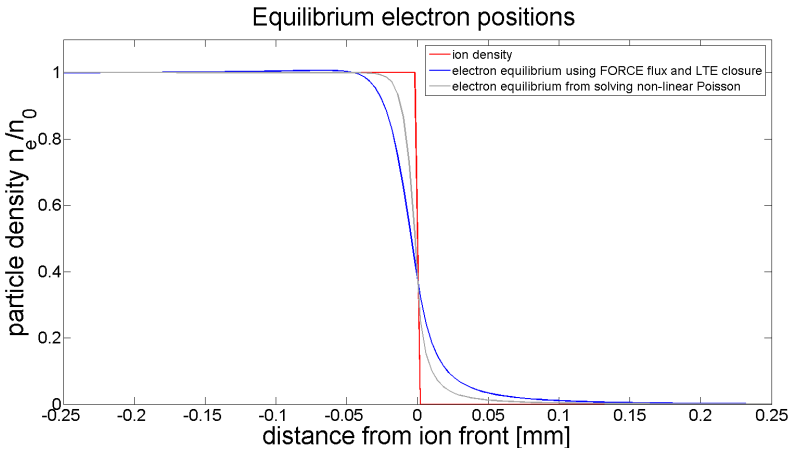

Figure 15: Electron distribution in equilibrium against a static ion distribution.

4.1.3 Electron equilibration

Although the model does not contain any explicit damping terms, the presence of the equation of state and the single-valuedness of the momentum imply instant thermalization of opposing streams. As a result, it is expected that high temperature electrons form an equilibrium configuration around a static positive charge distribution, as they do in a real experiment. Failure to do so would indicate a flaw in the model or its implementation. To test this, the familiar one-dimensional expansion scenario is set up, only with processing of the ion fluid disabled, resulting in a static ion distribution. As a result of the spatial dependence of the electron temperature, the distribution configuration is not expected to be exactly the same as in the isothermal case, but it should nonetheless be similar and free of undamped oscillations. The initial internal energy of the electron fluid is taken to be 1.0 MeV per electron, approximately twice the rest energy, so well in the relativistic regime. Figure 15 shows the resulting electron density distribution against a static ion distribution after 13.3 ps.

Note that in the case of fixed ions, the energy conservation is fairly bad and the electrons constantly gain energy resulting in a rising temperature that slowly increases the Debye radius and self-similarly broadens the electron density distribution around the ion border. The reason for this is that the numerical diffusion in the scheme moves electrons to an area of higher electrical potential without reducing the electron fluid's energy density. This numerical diffusion is also the reason the profile is too wide; the isothermal profile is implemented using a scheme that does not suffer from this diffusion. Regardless, this shows the numerical implementation of the LTE model produces stable results.

4.2 Maxwell solver

4.2.1 Dispersion relation of plane waves in vacuum and plasma

One well-understood plasma phenomenon that is easy to check against is the propagation (or lack thereof) of electromagnetic waves in a plasma. The dis-

persion relation for EM waves in plasmas is given by $\omega = \sqrt{\omega_p^2 + k^2c^2}$. In a simulation, two potential problems can occur. First, for time-domain Maxwell-solvers in general, due to the limited grid resolution, there is a limit to the spatial frequency that can be resolved. Most Maxwell solvers show a large deviation from the correct wave-frequency on waves with wave-vectors close to this limit. As explained in the methods section, the expected numerical dispersion relation is given by $\omega = \frac{2}{\Delta t} sin^{-1} \sqrt{S^2(sin^2(\frac{1}{2}k_x\Delta) + sin^2(\frac{1}{2}k_y\Delta) + sin^2(\frac{1}{2}k_z\Delta))}$. In the presence of a plasma, the limited time step size of a simulation may cause the plasma frequency to be impossible to resolve. Finally, in fluid simulations, mixing of electrons both as a result of the ensemble quantities being single-valued (opposing streams cancel out their momentum leaving just a high energy density, effectively thermalizing instantly) and as a result of numerical diffusion, can cause unphysical thermalization of the electrons. This results in the waves being damped more than they would be in a real plasma. An easy way to check for these effects is using Fourier analysis. First, the fields in the simulation area are initialized with white noise. This way, waves with all wave-vectors permitted by the simulation grid are present. A spatial Fourier transform (two or three-dimensional depending on the dimensionality of the simulation) of the transverse component of the electric field at each time step now gives the phase and amplitude of the waves for different wave vectors per time step. To find the frequency at which each wave oscillates, a second Fourier transform over time is performed on the spatial Fourier transforms, which shows an amplitude peak at this frequency, with a breadth depending on the duration of the oscillation. The location of these peaks can be directly compared to the analytical value.

The dispersion relation found from the field solver in vacuum is shown in figure 16. Similarly, the dispersion relation in the presence of an electron fluid is shown in figure 17

In the absence of an electron fluid, the Yee scheme is symmetrical in time, meaning it is fully energy conserving, resulting in the peaks in the Fourier transform being as sharp as the number of time steps allows. As expected for the Yee scheme, the dispersion relation is close to the analytical result for low spatial frequencies, but begins showing a large deviation from the correct result for higher spatial frequencies. As a result, it important that the simulation resolution picked such that the highest spatial frequency of travelling waves is low compared to the highest spatial frequency allowed by the grid. Given that without external EM radiation sources, the fastest oscillation in a plasma is given by the plasma frequency, the resolution required to negate the influence of the diffusion of the fluid is considerably larger than the resolution needed to keep the numerical dispersion of these waves low.

4.2.2 Convolutional Perfectly Matched Layer Boundaries

In the field solver, the simulation area boundary is terminated with a CPML, intended to dampen reflections of outgoing waves to approximate an infinitely large simulation area. To get an idea of how well such a boundary absorbs incoming waves, the following test was performed: in an initially empty two-dimensional simulation area, a current source is introduced slightly off-center. The field is measured in one of the two opposing corners. The same setup with a simulation area sufficiently large to prevent any reflections being measured is

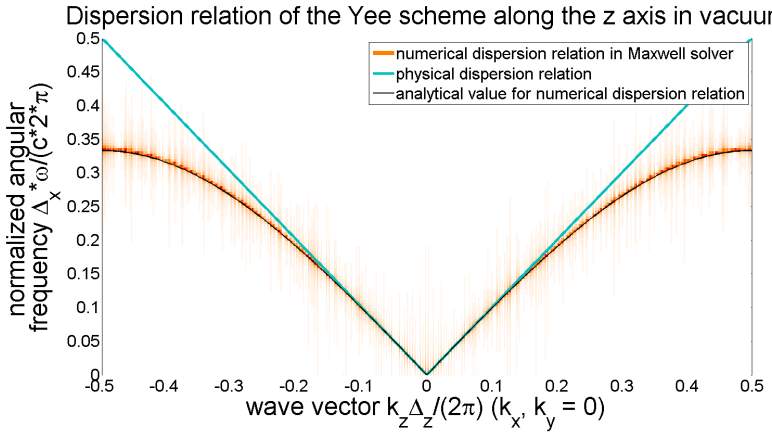

(a)

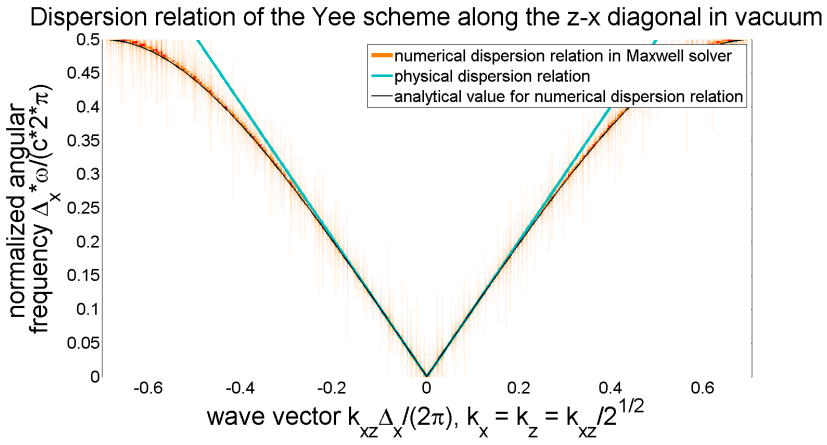

(b)

Figure 16: Dispersion relation diagram for plane waves in vacuum using the Yee solver implemented for the simulation. The numerical result is found after a spatial and temporal Fourier transform of the field in a two dimensional simulation initialized with noise in each field. As expected, the Yee scheme produces accurate results for low frequency electromagnetic waves, but shows a deviation from the correct frequency for shorter wavelengths.

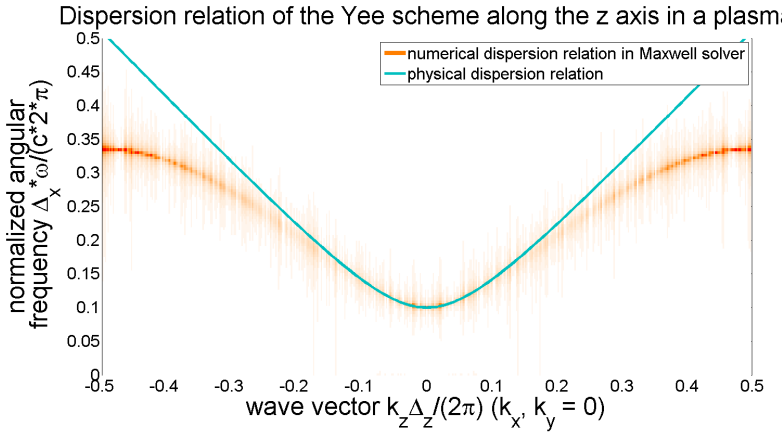

(a)

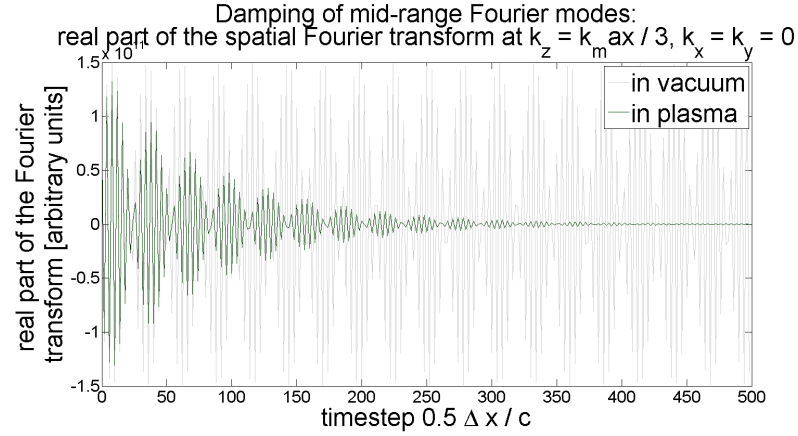

(b)

Figure 17: Dispersion relation for plane waves in the presence of an electron fluid, simulated with a 1st order central flux scheme (FORCE). The plasma frequency, the lowest frequency allowed in the plasma, is resolved correctly for the density (which was chosen such that the plasma frequency is the same as the frequency of waves at $k = k_max/10$). Whether or not the plasma frequency is resolved correctly depends on the time step taken. When the density becomes so high the plasma frequency comes close to the maximum frequency that can be resolved with this time step, the plasma frequency can no longer correctly be resolved. Unlike in vacuum, the sharpness of the peaks in the temporal Fourier transform is no longer limited by the duration of the simulation but by the damping of the modes as shown in the lower image. This is a result of the diffusion in the electron fluid. While the fluid scheme is energy conserving, opposing electron flows mix resulting in the electron fluid giving less energy back to the field. For very low spatial frequencies, this happens less quickly due to the opposing flows being further separated, while at very high frequencies, very little energy exchange happens between the field and the electrons, resulting in the damping being highest in the mid-range frequencies.

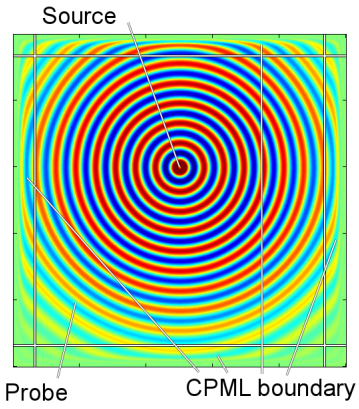

Figure 18: Cartoon of the test used to confirm the performance of the implementation of the CPML boundaries.

used as a control. The deviation from free propagation caused by reflections in the smaller simulation box is now calculated and normalized as $|\mathbf{E}_{probe,finite} - \mathbf{E}_{probe,free}|/E_{probe,max}$ where $E_{probe,max}$ is the biggest value measured at the probe during the simulation. The test is repeated for different frequencies to test for frequency dependence of the absorption.

The relative deviation from free-space propagation caused by the absorbing boundaries is shown in figure 19.

As can be seen in the graph, there is a strong frequency dependence in the absorption, however the low frequency reflection is still reduced by three orders of magnitude. Given that no strong radiating effects are expected in expansion simulations, it seems unlikely the finite simulation size would affect the final output much.

4.3 Full simulation

4.3.1 Energy conservation

The duration of the simulation is taken to be a nanosecond, which is of the order of the time needed for a particle beam moving at near light velocity to travel a meter. As can be seen in figure 20, the total deviation of the initial energy is about 1% of the total energy over that time.

To test the effect of dimensionality on energy conservation, a similar test is performed with a two dimensional simulation. The simulation area in this case represents the cross-section of a thin fully ionized wire, which expands cylindrically. Due to the processing power required the simulation has only been run for 0.14 ns. The change in the total energy of the electrons, protons and full simulation area is shown in figure 21 After the 0.14 ns period, there is a deviation of about 1.4% of the total energy.

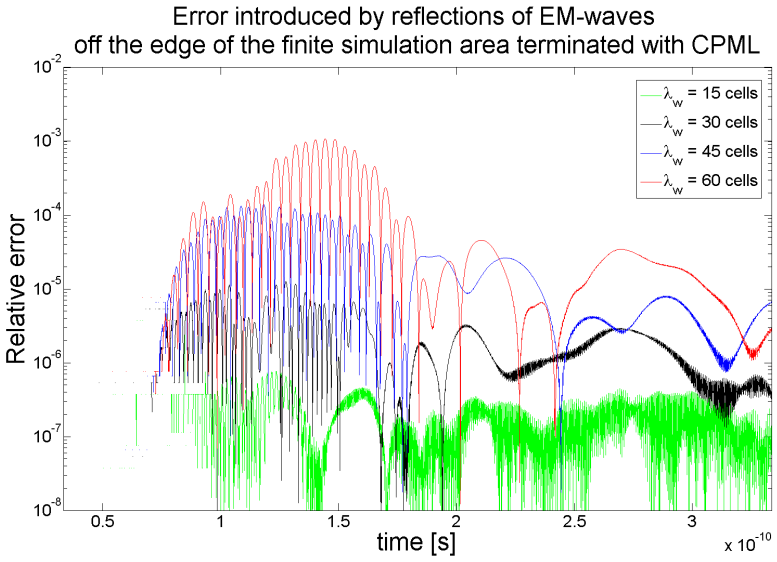

Figure 19: Relative size of reflections from the simulation border, attenuated by a convolutional perfectly matched layer (CPML). Although there is a strong frequency dependence for the absorption, even for low frequency radiation the reflections are attenuated by three orders of magnitude.

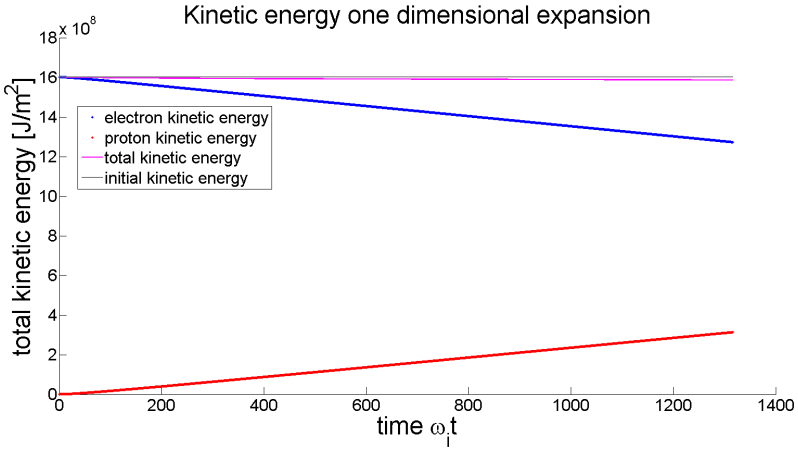

Figure 20: Total kinetic energy stored in each fluid species and the total in a 1 nanosecond one-dimensional simulation of an expanding infinitely thick foil, compensated for energy lost on the vacuum edge of the simulation and energy moved into the simulation on the foil side. The energy gained by the ions is compensated by energy lost by the electrons. Numerical errors result in a total deviation of about 1% over one nanosecond.

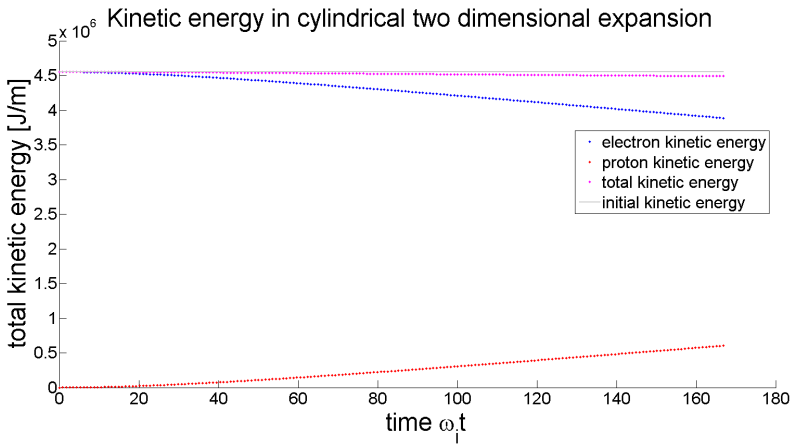

Figure 21: Total kinetic energy stored in each fluid species, as well as the sum of the energy, in a 0.12 ns two dimensional simulation of a fully ionized cylinder. Compared to the one-dimensional case the energy conservation is worse. As in the one-dimensional case, only an unphysical decrease of the energy is observed, as opposed to the increase seen in particle simulations.

4.3.2 Parallelization

For a large number of grid points in the simulation, processing required per time step will take more time than is practical using a single processor, and it becomes interesting to spread the workload out over multiple processes handled by different processors. Modern supercomputers have thousands of separate processor cores potentially allowing massive speedups of single simulations provided the calculations can be efficiently split between the different cores. The ability to run ever larger simulations not only depends on the amount of available processing power but also on the efficiency of this parallelization. Often, certain parts of a simulation cannot be parallelized and become bottlenecks limiting the maximum size. To test for the presence of such bottlenecks, two kinds of tests are common. The first is testing for strong scaling, where the total size of the simulation is held constant but the number of processors is increased. The other is testing for a simulation's weak scaling, where the amount of work per processor is held constant, so the total simulation size is increased along with the number of processes. In order to test a realistic scenario, the following tests were run on plasma expansion with the same parameters used in the physical results section. For the weak scaling test, in the one-dimensional case, each process is given 125 grid-points to process for 120,000 time steps. For the two-dimensional weak scaling test, each process is given 128x128 grid points for 4000 time steps. For the strong scaling test, in one dimension 32000 grid points are simulated for 120,000 time steps and in two dimensions 2048x2048 are simulated for 4000 time steps.

The fit is based on the assumption that the entire simulation is fully parallelizable (i.e. negligible serial code), but suffers from a relative overhead that scales with the number of processors to an unknown power a. This gives for N processors a speedup S of $S = T_1/T_N = N^{1-a}$. The optimal fits show a value

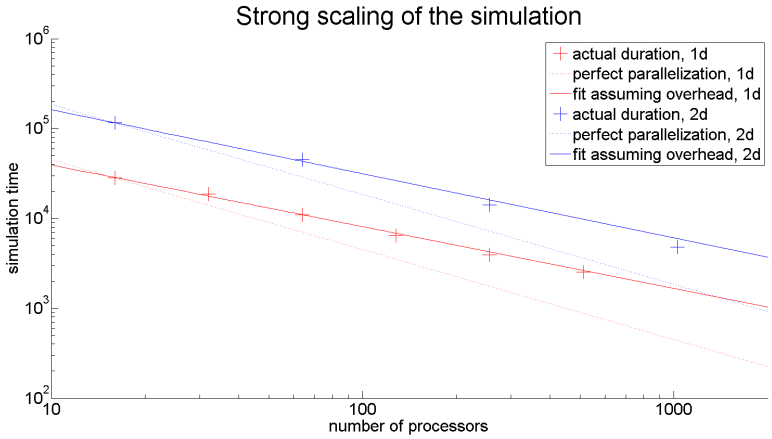

Figure 22: Strong scaling of a one and two-dimensional expansion. In the one-dimensional case, 32000 grid points are simulated. In the two-dimensional case, an area of 2048x2048 grid points is simulated. A strong deviation from the perfect case is visible.

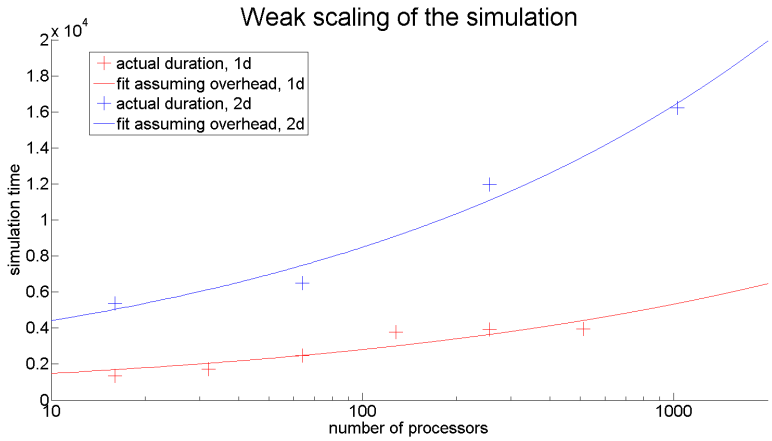

Figure 23: Weak scaling of a one and two-dimensional expansion. In the one-dimensional case, each core handles 125 grid points for 120,000 time steps. In the two-dimensional case, each core handles a 128x128 gridpoints square section for 4000 time steps. In the case of perfect parallelization, the simulation time would remain constant.

of a of approximately 0.28 to 0.29.

It is clearly visible that introducing more processors speeds up processing significantly as seen in figure 22, but also introduces a large overhead, as is visible from the deviation from the perfect parallelization and the increase in runtime in the weak scaling as shown in 23. Extrapolating the two-dimensional results shows that with a 10240x4096 grid, 2560 processes, a resolution of $\Delta = 2\mu m$ and a simulation time of $0.4ns$ the simulation would take a little over 5 days. Extrapolating to the three-dimensional case however, a 10240*4096*4096 simulation box at the same resolution with 8192 processors would take around 26 years.

4.4 Summary of test results

- The simulation output correctly converges for increasing resolution. A cell size of at most 2 μm is needed to get useful results.
- The flux limited central scheme does not work in the simulation and can therefore not be used.
- With LTE closure, the electrons correctly equilibrates against a static ion distribution, although the Debye length is overestimated.
- The Yee solver behaves as expected and when combined with the fluid solver correctly resolves the plasma frequency. The resolution required by the Yee solver is lower than for the fluid solver.
- The CPML boundary reduces boundary reflections of radiation by a factor of at least 1000.
- In plasma expansion, energy conservation of the full simulation scheme is good, resulting in a drop of only about 1% of the total energy over one nanosecond in one dimension.
- Parallelization works, but introduces an overhead. Currently full three-dimensional simulation is too slow for full-length simulations.

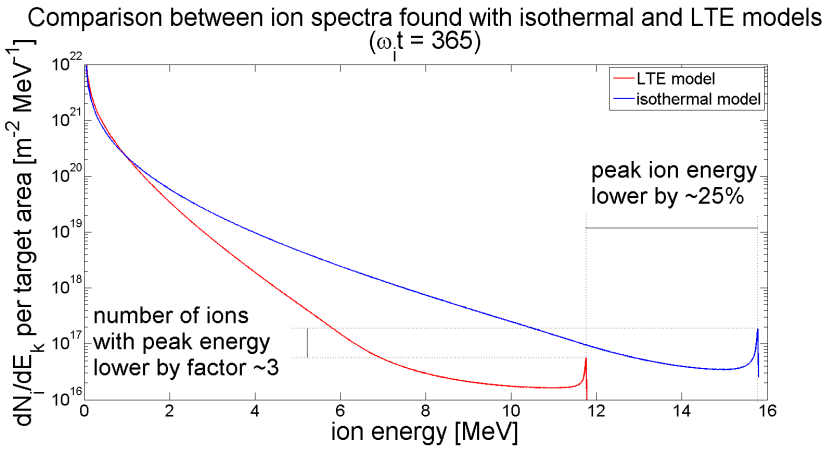

Figure 24: Ion spectrum at $\omega_i t = 365$ for the isothermal model and LTE model. At the same time after the expansion has started, the LTE model predicts a significantly lower peak ion energy, as well as a considerably smaller amount of ions accelerated to the peak energy.

5 Physics results

5.1 One-dimensional LTE expansion

Although the major advantage of LTE closure is the ease of generalization to higher dimensions, it produces different results from isothermal and adiabatic electron closures in one dimension. A comparison between a one dimensional expansion simulated using an implementation of the isothermal expansion model and the current implementation of the LTE model is presented in this section.

5.1.1 Simulation parameters

The one-dimensional LTE simulation was initialized with a very thick foil of protons and hot (1 MeV) electrons occupying 1/3rd of a $2cm$ long simulation area. The slab is thick enough that the rarefaction does not reach the simulation boundary. The resolution was selected based on the resolution tests in the tests section to be . The isothermal simulation has an identical initial state, but resolves the electron density and fields by approximately solving the nonlinear poisson equation with a constant internal energy of 1 MeV for the electrons.

5.1.2 Differences between isothermal model and LTE fluid models

Figure 24 shows a comparison of the ion spectrum predicted by both models. Two clear differences are visible. First, the maximum energy predicted is lower in the LTE model compared to the isothermal model. Second, the shape of the predicted spectrum is different with LTE closure and shows a plateau at high energies that is not present in the isothermal case. Both effects can be explained as follows.

- To confirm that the peak ion energy predicted by the LTE model is always lower, the evolution of that energy is shown in figure 26, from which it is

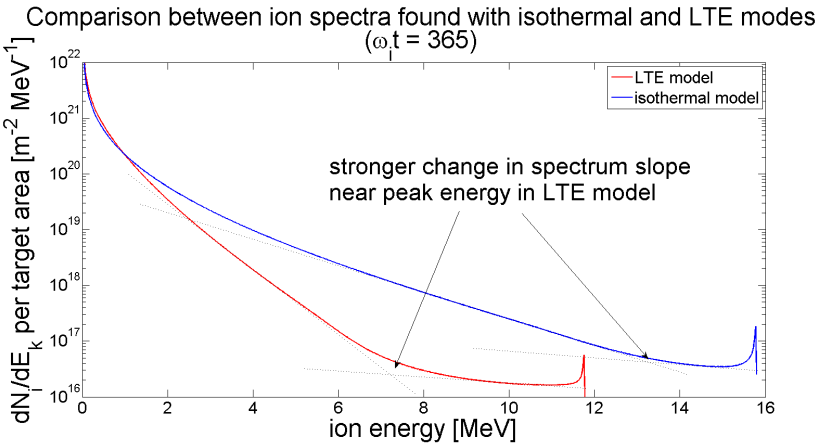

Figure 25: The same spectrum, now with the stronger deviation from the exponential spectrum found in the LTE closure highlighted.

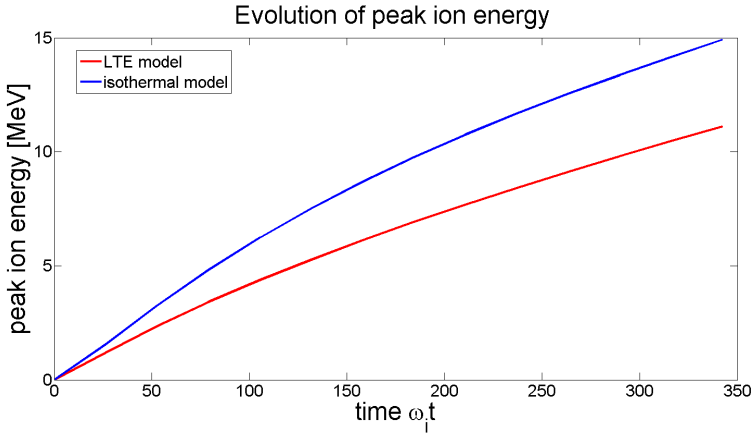

Figure 26: Peak ion energy in both models throughout the simulation, showing the highest ion energy found with the LTE model lags permanently behind the ion energy predicted by the isothermal model.

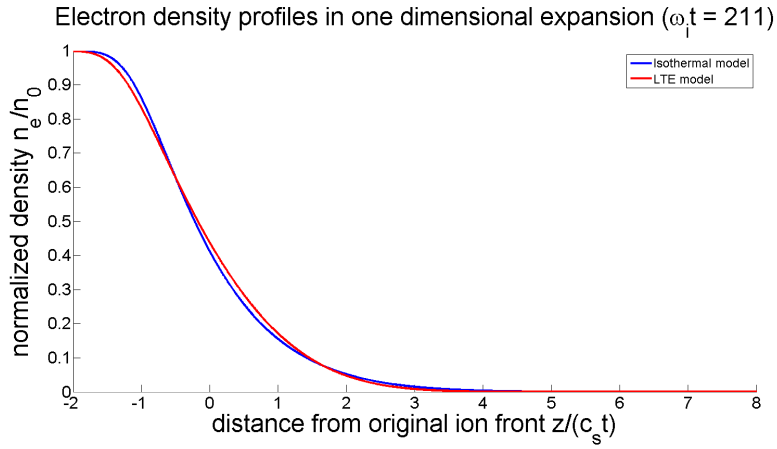

Figure 27: The electron density a while into the expansion. A clear difference is visible, with the LTE model showing a low electron density before the initial front, a higher density around the steepest part of the density gradient and a lower density further out. Note that the ion density is optically the same as the electron density because most electrons cannot go far away from the ions.

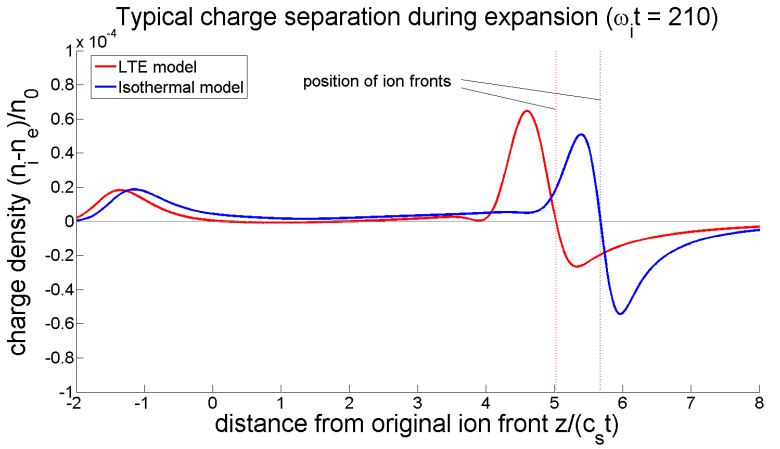

Figure 28: Charge separation during the one dimensional expansion. Both the charge separation in the isothermal model and the LTE model are shown. The LTE model shows less charge separation in the mid-region of the expansion.

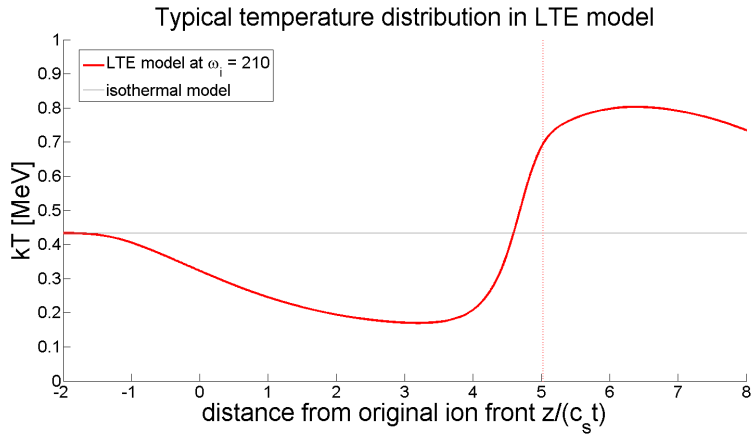

Figure 29: Plot showing the steep local dip in temperature predicted by the LTE model. This reduced temperature locally reduces the Debye length, which results in less charge separation and a smaller electric field.

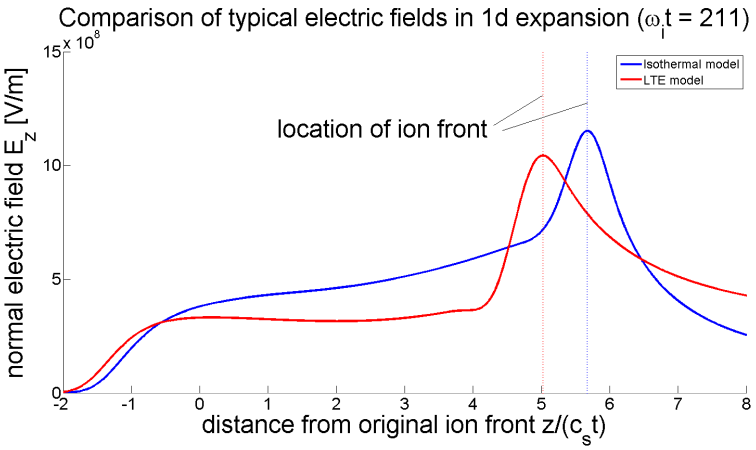

Figure 30: Normal electric field component throughout one dimensional expansion.

clear that the ion energy is indeed lower at every point in time.

- The maximum energy reachable by the ions is determined by the number of hot electrons and their initial temperature [10]. Since the two models are initialized the same, this must mean the electron density and or temperature develop differently. Figure 27 shows that, indeed, the electron density is different in the LTE case.

- The bulk of the electrons, however, cannot move further than the Debye length away from the ions, so the ion distribution looks very similar to the electron distribution. A better way to show how the electrons are positioned Figure 28 shows the difference in ion and electron density to show how the charges are separated. Since the diffusion from the fixed grid makes it impossible to pinpoint the exact ion front, we shall refer to the point where electrons outnumber the ions as the ion front. In the isothermal case there is a much larger negative charge peak directly outside the ion front, but the positive charge peak before the ion front is lower. Instead, the electrons outside of the ion front come from the entirety of the expansion region.

- A measure for how far the electrons can move away from the ions is the Debye length, which depends on the temperature of the plasma. Figure 30 shows the electron temperature in space during the expansion. This shows where the differences between the LTE model and isothermal model come from: the temperature in the LTE model shows a steep drop where the plasma has already expanded, while the isothermal model assumes the temperature is constant over space and time. The electrons from the cold region of the plasma remain relatively close to the ions at that point, resulting in a small charge separation, while the hot electrons at the ion front can move out further. The reason the temperature drops is because the plasma is expanding: the electrons that are pulled back by the ions lose some of their energy to the ions that are slowly accelerating outwards. Near the ion front, however the temperature is higher and the few electrons present here can move away further from the ions.

- The changed charge separation obviously also results in a different electric field, as shown in figure 30. The total number of electrons outside the ion front is lower in the LTE case due to the lowering temperature, so the peak normal electric field is also lower. This means the field in which the ions are accelerated is smaller, and the kinetic energy they receive is lower as well. Finally, from the fields, the change in the shape of the spectrum can also be explained. In the isothermal case the electrons come from everywhere in the expansion, and the field has a smooth buildup towards the peak. This accelerates the ions to many different energies evenly. In the LTE case, the field is more or less constant over a long spatial range and much lower compared to the peak, meaning a much bigger difference in the energy of ions that get accelerated at the peak and those that get accelerated in the field behind the peak.

5.2 Two-dimensional LTE expansion

5.2.1 Simulation parameters

In order to simulate a two-dimensional LTE expansion, the following initial conditions were set up. The two-dimensional simulation area is taken to be a cross-section of an infinitely wide thick foil with its normal in the positive z direction. It is heated with a Gaussian profile depending on x only: $\rho_e = (1 + ae^{-(x/w_0)^2})m_e c^2$ with $a = 2$ and $w_0 = 0.08cm$. The simulation area has a size of 2 cm in the z direction and 0.8 cm in the x direction and is assumed uniform in the y direction. The simulation area starts 0.66 cm into the foil. The density of the electrons and ions in the foil is initialized with a default low density of $n_0 = 10^24m^{-3}$. The far edge is initialized using an open boundary, and the other three with adiabatic boundaries. The proton fluid is simulated using the cold fluid closure, while the LTE closure is used for the electron fluid. The geometric setup and initial electron temperature distribution are shown in figure 31.

5.2.2 Differences between one and two-dimensional expansion

- Compared to the one-dimensional case, the peak ion energy is significantly lower, while the shape of the spectrum is more or less the same, as shown in figure 32.
- Figure 33 shows again the evolution of the peak ion energy. The reduced ion energy persists throughout the simulation. The acceleration of the ions is initially the same as in the one-dimensional case, which suggests the possibility that two-dimensional effects only become important once the the plasma has expanded a distance comparable to the size of the temperature profile.
- In figure 34 the distance the ion front has travelled is shown.
- The electron density is shown in 35. This shows that the expansion indeed happens both in length and width.
- Following the discussion of the one-dimensional case, the primary cause for the reduced maximum energy is the cooling down of the electron cloud before the ion front. Figure 36 shows the temperature distribution. Compared to the one-dimensional case, the relative decrease in temperature is considerably larger. Whereas in the one-dimensional case, the temperature dropped to about 1/3rd of the initial temperature, in two dimensions it drops to between 1/8th and 1/9th of the initial temperature.

5.3 Critical Analysis

5.3.1 Differences between the LTE model and existing models

Comparing the LTE fluid expansion with the isothermal model shows a clear difference: with the LTE closure, less energy is transferred to the hottest ions, even if slightly more energy is transferred to the colder ions. The primary reason for this is the drop in electron temperature before the ion front, a result of the expansion of the plasma as a whole. The finite potential created by the ions

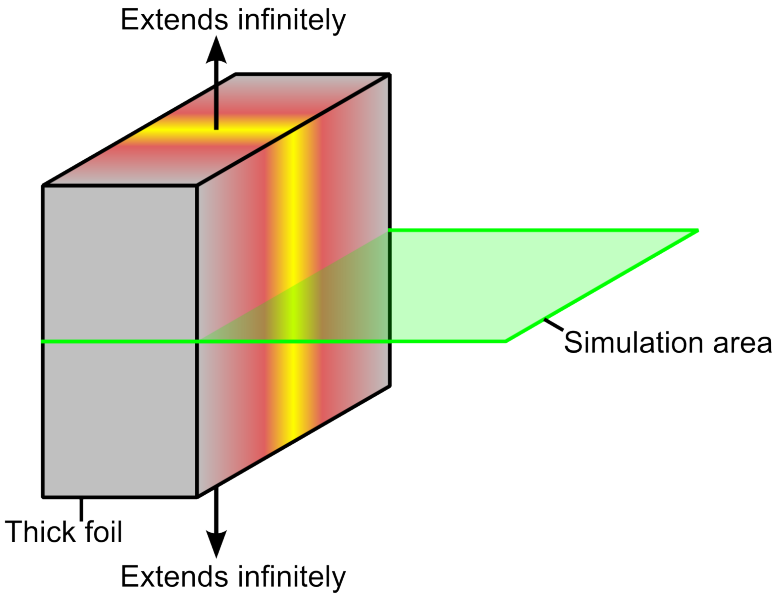

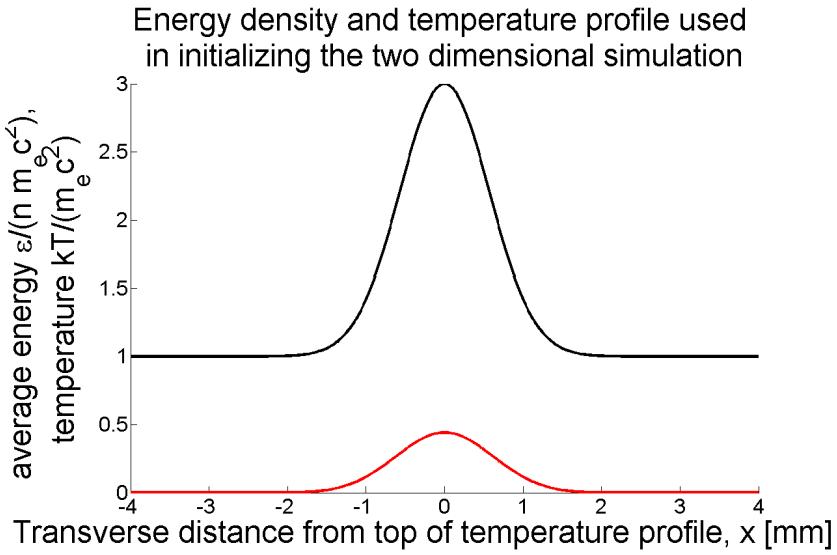

Figure 31: Set-up for the two-dimensional simulation. The simulation area is a cross-section of the foil, which is throughout heated with a Gaussian profile in one transverse coordinate, as by a laser with a line focus.

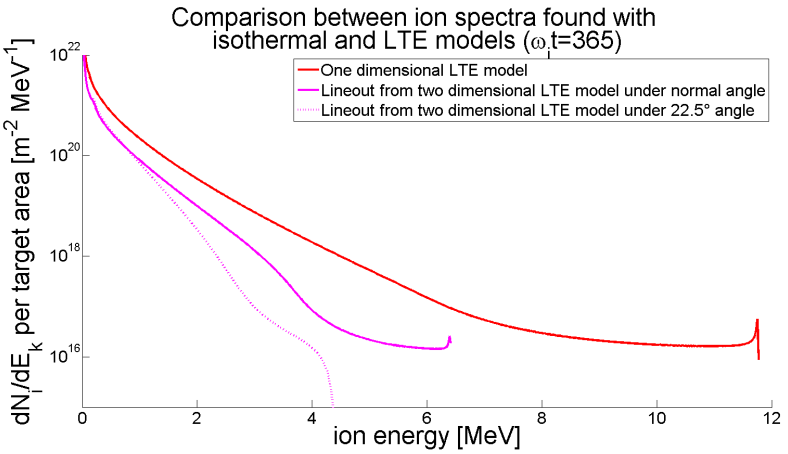

Figure 32: Comparison of the spectra in a single direction from the one-dimensional expansion and the two-dimensional expansion using the LTE model. The peak energy in the two-dimensional case is significantly lower.

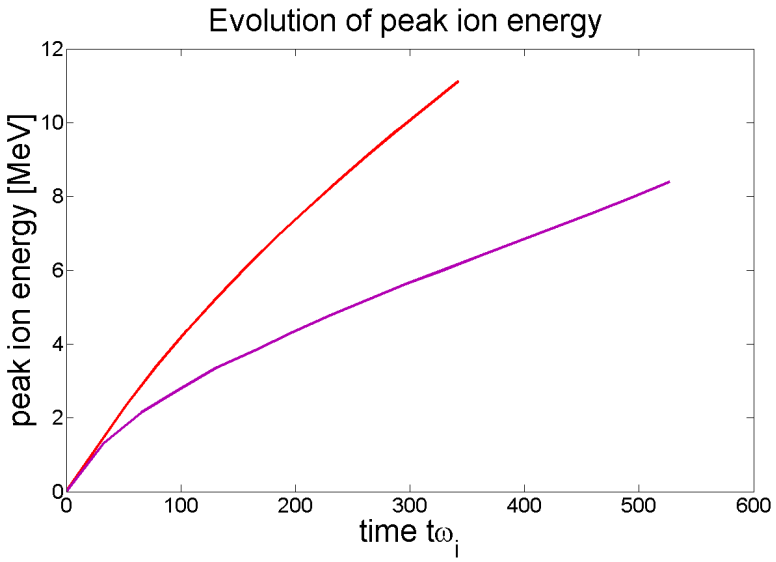

Figure 33: Peak ion energy in both models throughout the simulation for the one-dimensional and two-dimensional case using the LTE model, showing the much lower ion energy found when including two-dimensional effects.

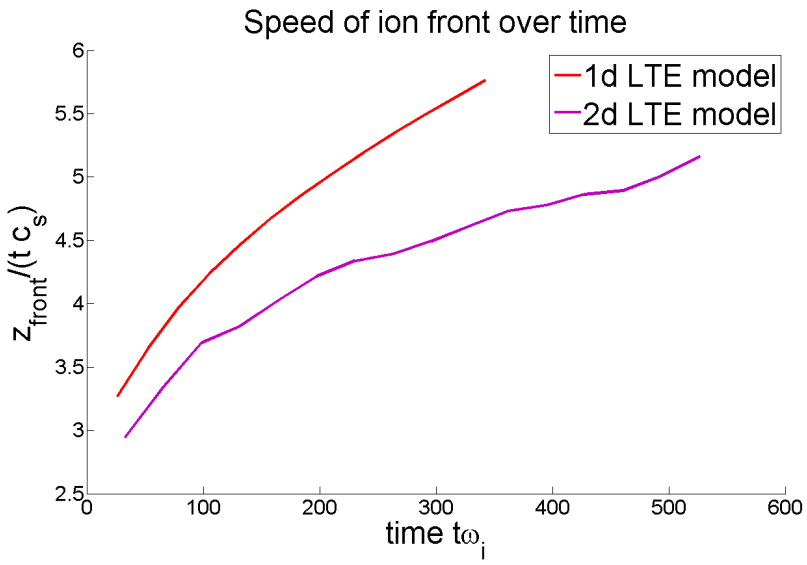

Figure 34: Comparison between the velocity of the ion front in the one and two-dimensional simulations. Due to propagating electromagnetic radiation in the two-dimensional simulation, the exact position of the ion front was harder to resolve. In both the one and two dimensional simulations, the ion front is taken to be the furthest position from the initial front where the number of ions exceeds the number of electrons. Note that due to the lower velocity found in the two-dimensional case, the plot has been extended to include the extra time where the ion front is still within the simulation area.

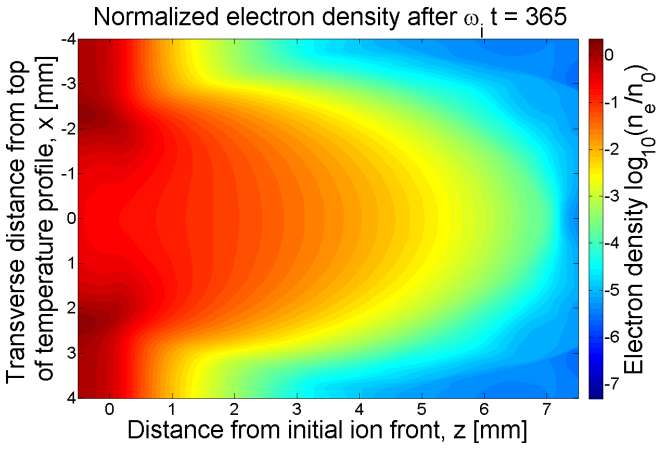

Figure 35: Electron density at $\omega_i t = 365$. The electron cloud expands more or less spherically from the initial position.

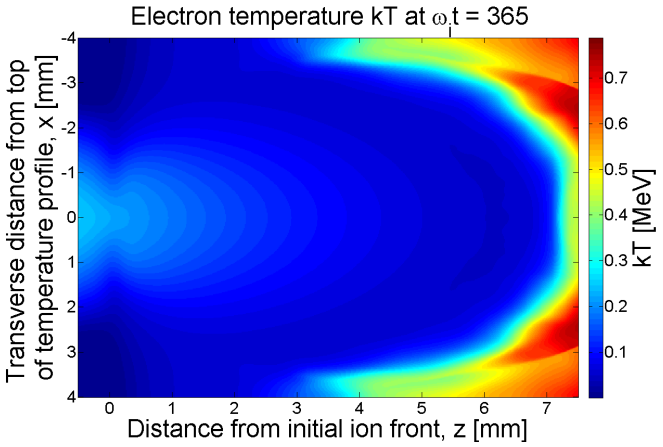

Figure 36: Temperature of the electron fluid at $\omega_i t = 365$. The lowest temperature is around 0.05 MeV, about one ninth of the original temperature in the foil. The fluctuations on the right are due to the propagating EM radiation.

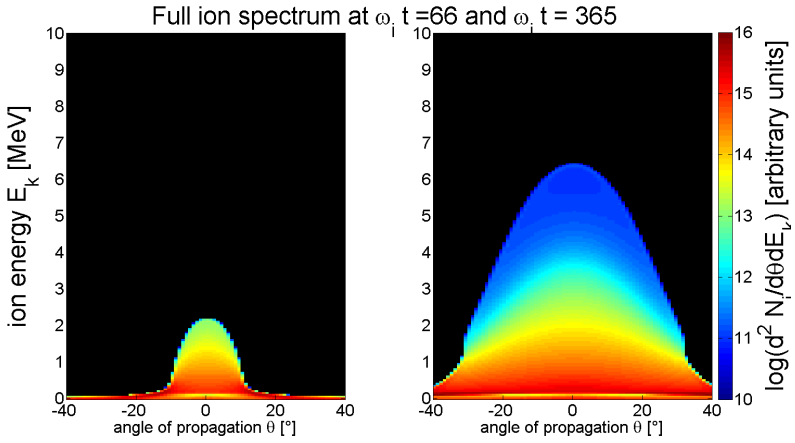

Figure 37: Ion energy spectrum from the two-dimensional simulation at $\omega_i t = 66$ and $\omega_i t = 365$. In the two-dimensional case, the angular spread of the expansion can also be resolved, something that was not possible in one dimension.

filters out the lower energy electrons away from the foil, leaving only the hottest electrons away from the ion front, resulting in a hotter temperature there. Newer one-dimensional models like the kinetic Mora model predict similar temperature distributions, but fail to generalize to multiple dimensions. The LTE expansion model can be upgraded to multiple dimensions easily. From the two-dimensional simulation, it can be observed that the cooling down effect seen in one dimension becomes bigger when including more dimensions.

5.3.2 Influence on experimental analysis

As mentioned in section 2.2.4, the energy passed to each electron from an impacting laser pulse on a solid target scales with the peak intensity of the pulse. In the plasma-vacuum expansion happening in the final stage of TNSA, the electron temperature determines, among other factors, the size and amplitude of the fields in which the ions are accelerated, and therefore the final spectrum of the ions. This effectively means that a model that predicts a lower peak kinetic energy, also predicts that in order to reach a certain ion energy from the TNSA process, a higher laser intensity is needed. In the case that the expansion model is only used to investigate what happens after the initial stage of the expansion, it only predicts what happens with the energy stored in the electrons at that point. Depending on how much energy that is, this might still influence the spectrum of the ions. Using an existing one-dimensional model would predict that the electrons still give off energy to the ions in the normal direction, however from PIC simulations it is known that the electron distribution quickly thermalizes and becomes isotropic [27], which means that part of the energy cannot be passed to the ions in the normal direction. As a result, a two-dimensional model as discussed in this thesis will predict less energy passed to the fastest ions and a bigger angular spread.

6 Conclusions

We have developed, implemented and tested a fully relativistic fluid model for plasma-vacuum expansion and applied it to the case of the expansion phase of target-normal sheath acceleration. The fluid model allows long timescale simulation of plasma expansion in any dimensionality, limited only by available processing power and implementation efficiency. Simulations based on the model are easily parallelized due to the model being fully causal; only local information is needed to advance the state of the system. Compared to existing models for plasma expansion, the LTE fluid model has much less strict assumptions. There is no need to assume a global temperature for the electrons, as the temperature of the electrons is determined locally based on the local energy density. It does not need the assumption of an energy source or a an a-priori expansion time. Finally no relativity-violating solution of the non-linear Poisson equation is needed. Currently well known models of plasma-expansion use at least one of these assumptions.

This model allows to bridge the gap between what happens between the initial stages of the ion acceleration and the actual measurements done in experiments. In addition, it is now possible to investigate things like the electron spectrum and spatial distribution and the development of the angular dependence of the ion spectrum over time.

We have applied the LTE model to a simple thick-foil TNSA case in one and two dimensions and compared the results with the commonly used one-dimensional isothermal expansion model. The results show clear differences with the isothermal fluid model. Less energy is transferred to the fastest ions, and the final ion spectrum is less regular. These differences are much bigger in the two-dimensional case, indicating that multi-dimensional effects may play an important role in determining the final spectra. It is this ion spectrum that is actually measured in TNSA experiments, and the models that are currently used to compare experimental results with do not capture these multidimensional effects. As a general matter, existing plasma expansion models overestimate the maximum energy of the ions that can be reached with a certain laser intensity and pulse energy. The prediction of an overall lower energy in the two-dimensional expansion case suggests that including multidimensional effects is a step in the right direction.

7 Outlook

The LTE model can be applied to many more cases than discussed in this thesis.

7.1 Larger simulations

7.1.1 Multi-electron species expansion

As seen in the two-dimensional simulation, the single-valuedness of the momentum density results in a pile-up of electrons next to a locally heated spot instead of the cold counter-current that would be expected and can be observed in kinetic simulations. One possible solution to this problem could be to create two electron fluids, a hot and a cold one. This way, the hot electrons can expand almost freely into the region of the cold electrons, because the cold electrons are accelerated in the opposite direction. This will allow comparing the different peak ion energies reached with wide foils (where the hot electrons are not confined) and small targets (where the hot electrons are pulled back at the edge of the target) Similarly this allows research into whether it is more efficient to increase the maximum electron temperature (higher pulse intensity) versus the number of hot electrons (higher pulse energy), since these can be treated with separate fluids while still not requiring any global temperatures.

7.1.2 Three dimensional expansion

In order to investigate three dimensional effects, two approaches can be used. In the case that the expansion is expected to be radially symmetric, the differential equations can be modified to represent a radially symmetric system in what is colloquially referred to as 2.5-D; this would not be significantly more computationally expensive than the two-dimensional simulation shown in this thesis, but would not allow simulating asymmetric expansions. The other approach is using the simulation code developed for this thesis as is to simulate a three dimensional expansion in Cartesian coordinates. Although low resolution tests have shown that the simulation can run in three dimensions, to run at a useful resolution this would require an equal grid size in the y direction as in the x direction. For the simulation as shown in section 5.2 this would mean an increase in required operations of a factor 4000. Going by the parallelization efficiency found in section 4.3.2 meaning the simulation would take many years to complete assuming the same processing power used with the same parallelization efficiency. For a fully three dimensional simulation, some speed-up is needed over the current scheme.

7.1.3 Upwind schemes and higher order schemes

The high resolution for the simulations discussed in this thesis was chosen to compensate for the numerical diffusion caused by using a scheme using a linear centralized flux approximation, which requires a large numerical diffusion term to stabilize the simulation. To simulate larger or 3-D systems, being able to use a lower grid resolution would be highly beneficial. As shown, reaching the higher order accuracy needed for lower grid spacing using centralized schemes proved problematic in this model. A possible solution would be switching from

centralized schemes to upwind methods, which use knowledge of the conservation equations themselves to achieve considerably lower numerical diffusion even in first-order schemes [8]. Because these schemes normally require solving an eigenvalue problem on each cell, they tend to be highly complex and computationally intensive, however in higher dimensions a reduction in resolution becomes more worthwhile (halving the resolution in all three dimensions allows simulating the same system 16 times faster; there are 8 times less fluid elements and the time steps can be twice as large by the CFL condition).

Similarly to further relax the CFL condition and provide better damping of outgoing waves at the simulation borders, the Maxwell solver could be replaced with the directional splitting scheme, which is an upwind Maxwell solver. In addition to having reduced numerical dispersion, the directional splitting scheme has a less strict CFL condition ($c\Delta t \leq \Delta x$, as opposed to $c\Delta t \leq \sqrt{1/D}\Delta x$ where D is the dimensionality). In three dimensions, this allows a further reduction of the number of time steps needed by a factor 1.73. In addition, this would allow simulation of open boundaries that don't suffer from the reflections seen in the two-dimensional simulation results in this thesis.

7.2 Hybrid simulations

Like every model, fluid models have a validity range in which they can be used. In particular, fluids cannot automatically model wave breaking (faster particles overtaking slower particles) or counterpropagating flows, making fluid simulation normally unsuitable for intense laser-matter interaction. On the other end of the spectrum, the fluid model loses its validity if the particle density becomes so low that talking about local averages no longer makes sense. For this purpose it may be necessary to combine fluid schemes with other methods of simulation.

7.2.1 Input from PIC to fluid

If the entirety of the TNSA process has to be simulated, this involves laser-matter interaction in the electron heating phase and obviously wavebreaking during the electron transport phase. (see section 2.2.3) As such a fluid simulation will not be able to simulate the first two stages, which require a full kinetic approach. This can be done using the Particle In Cell simulation. It is possible to take the output from a PIC simulation at the point where the electron fluid is more or less equilibrated at least locally, use this to calculate the averages (particle density, momentum density and energy density) for cubic cells, calculate the electric field in such a way that it is consistent with the charge density and and use this as the initial state of a fluid simulation.

7.2.2 Test-particles in fluid simulation

Although thus far, it was assumed that there is only a single ion species present in the expanding plasma, it has been known for some time that the highest energy ions measured in real TNSA experiments are not ions of the target material, but protons from the back surface of the target. These protons come from water vapor contamination that is normally present. Because these protons are the ions with the highest charge/mass ratio present, they stay ahead of the

other ions during acceleration, where the electric field is the largest. Because the number of protons is lower than the other ions, a situation could arise where their density becomes too low to describe as a fluid. A possible solution would be to include the protons as particles that are driven by the field but do not deposit their currents. Similarly, it would be worthwhile to include a number of test-electrons to observe electron trajectories and verify whether the fluid model provides a correct electron density evolution and check for possible wave breaking. Finally, in order to model wavebreaking of the ions, it would be possible to simulate the electrons using LTE closure and the ions using the PIC algorithm in a similar way as is currently done in most one dimensional models.

7.2.3 Multi-Physics / multi-scale simulations

In order to develop a truly versatile code that can tackle TNSA and possibly other complicated plasma processes, it would be worth trying to resolve different parts of the same simulation area using different models. In the case of TNSA for example, the foil itself could be simulated using a hybrid PIC code, which could include some form of molecular dynamics to simulate the untouched foil, code to calculate ionization rates at the front of the foil to simulate the formation of the plasma, classic PIC algorithms to handle the forming plasma and electron transport inside the foil, fluid schemes to calculate the expanding mid-outer reaches of the plasma and simple particle tracking to calculate the particles emitted from the plasma.

8 References

[1] Satish Balay, Jed Brown, Kris Buschelman, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc Web page, 2013. http://www.mcs.anl.gov/petsc.

[2] Guillermo Chacón-Acosta, Leonardo Dagdug, and Hugo A. Morales-Técotl. Manifestly covariant jüttner distribution and equipartition theorem. *Phys. Rev. E*, 81:021126, Feb 2010.

[3] F.F. Chen. *Introduction to Plasma Physics*. Springer, 1995.

[4] G.Q. Chen and E.F. Toro. Centered difference schemes for nonlinear hyperbolic equations. *Journal of Hyperbolic Differential Equations*, 01(03):531–566, 2004.

[5] Stephen P. Hatchett, Curtis G. Brown, Thomas E. Cowan, Eugene A. Henry, Joy S. Johnson, Michael H. Key, Jeffrey A. Koch, A. Bruce Langdon, Barbara F. Lasinski, Richard W. Lee, Andrew J. Mackinnon, Deanna M. Pennington, Michael D. Perry, Thomas W. Phillips, Markus Roth, T. Craig Sangster, Mike S. Singh, Richard A. Snavely, Mark A. Stoyer, Scott C. Wilks, and Kazuhito Yasuike. Electron, photon, and ion beams from the relativistic interaction of petawatt laser pulses with solid targets. *Physics of Plasmas*, 7(5):2076–2082, 2000.

[6] R.W Hockney. Measurements of collision and heating times in a two-dimensional thermal computer plasma. *Journal of Computational Physics*, 8(1):19 – 44, 1971.

[7] R.W. Hockney and J.W. Eastwood. *Computer Simulation Using Particles*. Taylor & Francis, 2010.

[8] J.M. Ibez and J.M. Mart. Riemann solvers in relativistic astrophysics. *Journal of Computational and Applied Mathematics*, 109(12):173 – 211, 1999.

[9] Thomas Kiefer, Theodor Schlegel, and Malte C. Kaluza. Plasma expansion into vacuum assuming a steplike electron energy distribution. *Phys. Rev. E*, 87:043110, Apr 2013.

[10] T. Kluge. *Enhanced Laser Ion Acceleration from Solids*. PhD thesis, T.U. Dresden, 2012.

[11] Mark J.H. Luttikhof. *Theoretical investigation of external injection schemes for laser wakefield acceleration*. Ipskamp Drukkers, 2010.

[12] Andrea Macchi, Marco Borghesi, and Matteo Passoni. Ion acceleration by superintense laser-plasma interaction. *Rev. Mod. Phys.*, 85:751–793, May 2013.

[13] P. Mora. Plasma expansion into a vacuum. *Phys. Rev. Lett.*, 90:185002, May 2003.

[14] P. Mora and T. Grismayer. Rarefaction acceleration and kinetic effects in thin-foil expansion into a vacuum. *Phys. Rev. Lett.*, 102:145001, Apr 2009.

[15] P. Mora and R. Pellat. Self-similar expansion of a plasma into a vacuum. *Physics of Fluids*, 22(12):2300–2304, 1979.

[16] K. W. Morton and D. F. Mayers. *Numerical Solution of Partial Differential Equations*. Cambridge University Press, second edition, 2005.

[17] Alexander Piel. *Plasma Physics: An Introduction to Laboratory, Space, and Fusion Plasmas*. Springer, 2010.

[18] J. Alan Roden and Stephen D. Gedney. Convolution pml (cpml): An efficient fdtd implementation of the cfspml for arbitrary media. *Microwave and Optical Technology Letters*, 27(5):334–339, 2000.

[19] J. Schreiber, F. Bell, F. Grüner, U. Schramm, M. Geissler, M. Schnürer, S. Ter-Avetisyan, B. M. Hegelich, J. Cobble, E. Brambrink, J. Fuchs, P. Audebert, and D. Habs. Analytical model for ion acceleration by high-intensity laser pulses. *Phys. Rev. Lett.*, 97:045005, Jul 2006.

[20] 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, 2008.

[21] R. A. Snavely, M. H. Key, S. P. Hatchett, T. E. Cowan, M. Roth, T. W. Phillips, M. A. Stoyer, E. A. Henry, T. C. Sangster, M. S. Singh, S. C. Wilks, A. MacKinnon, A. Offenberger, D. M. Pennington, K. Yasuike, A. B. Langdon, B. F. Lasinski, J. Johnson, M. D. Perry, and E. M. Campbell. Intense high-energy proton beams from petawatt-laser irradiation of solids. *Phys. Rev. Lett.*, 85:2945–2948, Oct 2000.

[22] A. Taflove and S.C. Hagness. *Computational Electrodynamics: The Finite-Difference Time-Domain Method*. The Artech House antenna and propagation library. Artech House, Incorporated, 2005.

[23] EF Toro and SJ Billett. Centred tvd schemes for hyperbolic conservation laws. *IMA Journal of Numerical Analysis*, 20(1):47–79, 2000.

[24] M. A. True, James R. Albritton, and E. A. Williams. Fast ion production by suprathermal electrons in laser fusion plasmas. *Physics of Fluids*, 24(10):1885–1893, 1981.

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

[26] K Zeil, S D Kraft, S Bock, M Bussmann, T E Cowan, T Kluge, J Metzkes, T Richter, R Sauerbrey, and U Schramm. The scaling of proton energies in ultrashort pulse laser plasma acceleration. *New Journal of Physics*, 12(4):045015, 2010.

[27] K. Zeil, J. Metzkes, T. Kluge, M. Bussmann, T. E. Cowan, S. D. Kraft, R. Sauerbrey, and U. Schramm. Direct observation of prompt pre-thermal laser ion sheath acceleration. *Nature Communications*, 3:874+, June 2012.

Acknowledgements

This thesis was written at the HZDR in Dresden, Germany. It has been a long and occasionally stressful time, but more than anything it has been an enjoyable time, and I would like to thank everybody that has made it possible. First and foremost, I would like to thank my direct supervisor, Michael Bussmann. When I started this work he introduced me to plasma physics faster than I would have thought possible, but he has taught me an incredible amount more in the time that followed. I would like to thank Fred van Goor and Klaus Boller for their patience and continuing support from Enschede. Also, I would like to thank Bert Bastiaens for bringing me into contact with the people at the HZDR in the first place. Thanks go out to Thomas Kluge, for his help with PIC simulations, and Alexander Debus for his help with Mathematica among other things. A big thanks to the examination committee, especially Marc Dhallé as the external member, for reading this text. Many thanks to my fellow students at the theoreticists group at the HZDR. Working, talking and laughing with all of you has made the time spent working on this thesis as great as it was. Finally, I would like to thank my parents; without their continued and active support, none of this would have even remotely been possible.