CHALMERS

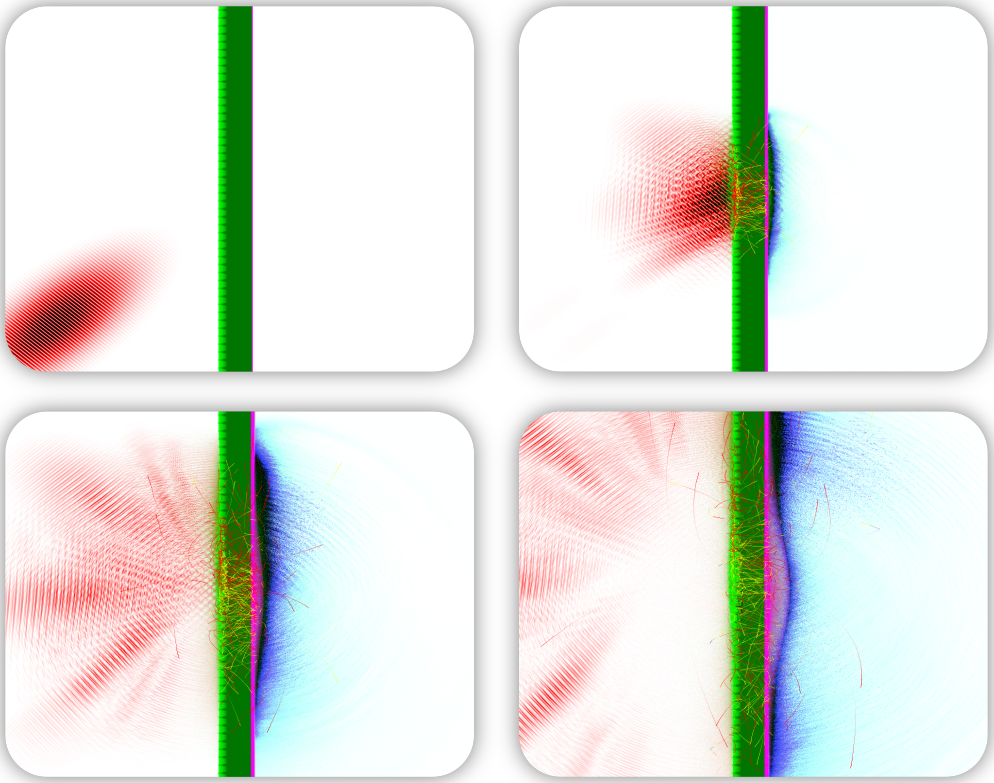

Particle-in-cell simulations of intense laser-matter interactions, with a focus on light ion acceleration

Master's thesis in Applied Physics

JOEL MAGNUSSON

Department of Applied Physics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2015

MASTER'S THESIS IN APPLIED PHYSICS

Particle-in-cell simulations of intense laser-matter interactions, with a focus on light ion acceleration

JOEL MAGNUSSON

Department of Applied Physics
Division of Condensed Matter Theory
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2015

Particle-in-cell simulations of intense laser-matter interactions, with a focus on light ion acceleration
JOEL MAGNUSSON

Master's thesis 2015:01
ISSN 1652-8557
Department of Applied Physics
Division of Condensed Matter Theory
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone: +46 (0)31-772 1000

Cover:
Target normal sheath acceleration using an ultra-intense laser pulse and a target with nano-sized cones on the front surface. The sub-figures are taken at different moments in time and show the electron density (green), proton density (magenta), absolute value of the electric field (red) and x-component of the electric field (blue). The lines show electron trajectories with colors to indicate particle energies (blue-yellow-red).

Chalmers Reproservice
Gothenburg, Sweden 2015

Particle-in-cell simulations of intense laser-matter interactions, with a focus on light ion acceleration
Master's thesis in Applied Physics
JOEL MAGNUSSON
Department of Applied Physics
Division of Condensed Matter Theory
Chalmers University of Technology

Abstract

In conventional accelerators the risk of electrical breakdown limits the electric fields used to accelerate charged particles to the order of 100 MV/m. As a result, the accelerators are required to be very long in order to achieve particle beams of very high energy. Owing to this, and with the advancement of lasers of ever increasing intensity and power, laser-driven plasma-based acceleration has become a field of intense study. This is mainly motivated by the fact that such accelerators can sustain acceleration gradients several orders of magnitude larger than conventional accelerators (fields of several TV/m have been measured), thus allowing the accelerator to be made correspondingly more compact.

In this work we focus on light ion acceleration in the TNSA regime (Target Normal Sheath Acceleration). In doing so, we also investigate unphysical numerical effects in the code Picador, a purpose-built numerical tool specifically designed for large-scale 3D particle-in-cell plasma simulations and with unique capabilities of investigating ultra-intense laser-plasma interactions. Picador is further developed with tools to more easily identify new numerical issues and also for suppressing found numerical effects on TNSA simulation results.

The issue of artificially charging the simulation boundary is shown to have little to no effect on TNSA simulations using Picador, given that the simulation region is sufficiently large. Furthermore, this size coincides well with the required size for acceleration saturation to occur and is shown to be feasible for 2D TNSA simulations with regards to computational costs. Furthermore, a number of numerical issues were identified, and eventually resolved, much through the use of a new module in Picador that simply tracks the total simulation energy.

Alternative target designs for improving light ion beam energy and collimation are simulated and analysed. Results show that the beam energy can be greatly improved with nano-sized cones on the front target surface. It is further shown that a similar effect can be obtained by restricting the hot electrons from being transported away in the transverse (to the target normal) directions. Furthermore, simulations using nano-sized cones on the rear target surface, in an attempt at increasing the strength of the charge-separation field, have been carried out. However, it is shown that this completely ruins the collimation of the beam and without improving the beam energy, the structures presumably being too small to affect the charge-separation field for a sufficient amount of time.

Acknowledgements

I would first like to express my deepest gratitude towards Arkady Gonoskov, who always puts his own things aside to help me out whenever I pester him with some new problem of mine. His guidance and input has truly been of great value to me throughout this entire work, helping me out with anything from proofreading to code debugging, from data interpretation to simulation suggestions and anything in between. He has not only been a terrific supervisor, but also a great friend.

I would also like to acknowledge Sergey Bastrakov, at Lobachevsky State University of Nizhni Novgorod, Russia, for his quick responses to my many questions regarding both coding in general and Picador. He has tirelessly, and at any time of the day, helped me sort out many and most of the issues that I have faced when performing my simulations and I do not believe I would have gotten as far as I did, without his assistance.

Furthermore, it is with a huge amount of gratefulness that I wish to acknowledge Mattias Marklund who set me up for this project. He is an inspiration to me and his supervision has truly helped in keeping my mind happily distracted from the problems at hand. I always know that I can go to his office for encouragement, guidance or just for a fun discussion[1].

I also wish to thank everyone, especially those not already mentioned, in the A-team here at Chalmers – Anton, Arkady, Chris, Greger, Erik, Felix, Mattias and Tom – for their warm welcoming and for the stimulating, not to mention fun, environment they provide on a daily basis. I am very grateful to have been given this opportunity to both work with and get to know you all.

Finally, I should also take this moment to thank my family for being ever understanding and supportive throughout my life, even though that often means that I have less time to spend with them than I would have wanted. With their support I have always known that I may pursue whatever path I find interests me the most.

Joel Magnusson, Gothenburg, August 25, 2015

[1] ..perhaps mostly the last

Contents

Chapter 1

Introduction

The energetic particle beams produced in conventional accelerators are today of great importance in a wide variety of areas of science, technology and medicine. The main use can be found in a range of different imaging applications, but the beams can also be used to study cosmic radiation damage both to spacecraft components and living organisms, to study wear of mechanical parts using so-called thin layer activation and for treating cancer tumors through radiation therapy. Conventional accelerators are however limited to electric fields used to accelerate the charged particles on the order of 100 MV/m, due to the risk of electrical breakdown. As an effect, very large and expensive accelerators have to be built in order to achieve particle energies sufficient for applications.

Laser-driven plasma-based accelerators on the other hand, are not limited by this and acceleration gradients several orders of magnitude larger than for conventional accelerators can therefore be achieved. As a consequence, laser-plasma accelerators can be made much more compact with the possibility of simultaneously reducing costs. This coupled with the rapid evolution of intense laser systems, see Figure 1.1, has made the field of ultra-intense laser matter interactions a hot research area in recent years.

1.1 Current status

Since the advent of the first functioning laser in 1960, by Maiman,[1] the capabilities of lasers have continuously been improved, to the extent of providing electromagnetic pulses of ultra-high intensities and ultra-short duration.[2] The research on the interaction between strong laser fields and plasmas began in the 1970s and was mainly focused on the topic of laser-induced nuclear fusion. Many scientific results of interest for particle acceleration were reported and of particular interest for ion acceleration is for instance the effect of resonance absorption in which the laser energy goes into so-called hot electrons.[3,4,5] The energy of the hot electrons is then, at least partially, converted into energetic ions via charge-separation fields.[6] The study of these effects where of importance in the inertial confinement fusion-context as the hot electrons causes unwanted preheating of the compressed plasma while, at the same time, the energetic ions steal a significant fraction of the laser energy.

It was not until ultra-short and ultra-high peak power lasers appeared in the 1990s that the interest in laser-plasma interactions grew as a means for ion acceleration, even though plasma-based ion acceleration had been proposed as early as 1956 by Veksler.[7] This rapid development in laser technology, Figure 1.1, is to a large extent owing to the invention of the Chirped Pulse Amplification (CPA) technique in 1985,[8] in which the pulse is stretched both temporally and spectrally prior to being amplified and after which it gets re-compressed, coupled with the discovery of the titanium-sapphire lasing medium in

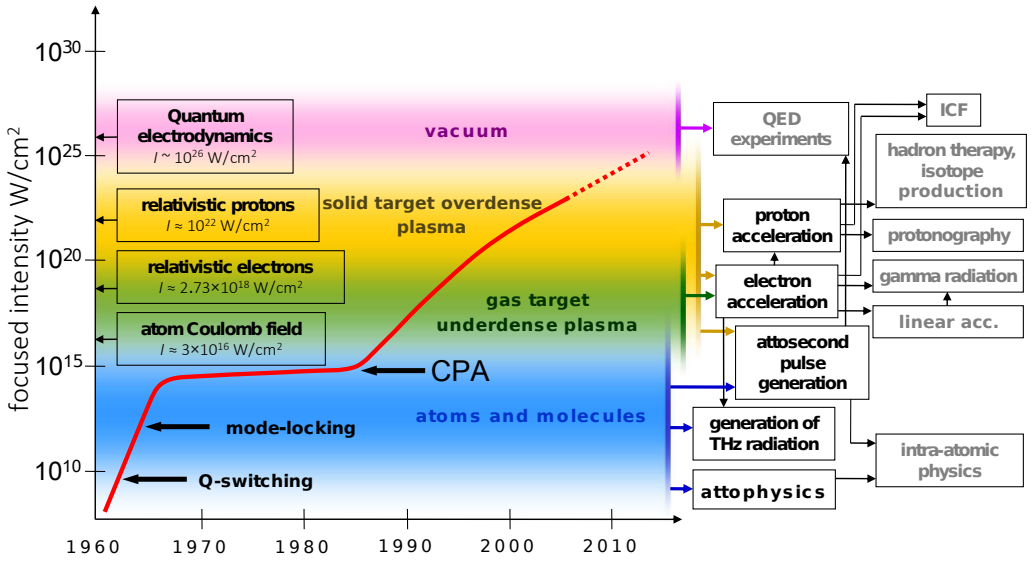

Figure 1.1: *Overview of the historical advancement of high intensity laser systems. The red curve represents the available laser intensity since the invention of the laser and the colored regions indicate the basic systems in which the laser-matter interactions exhibits nonlinear behavior. A selection of applications are shown to the far right in the figure. The figure is reprinted from [15] with permission of the author.*

1982.[9,10]

There are today a large number of high-power laser facilities around the world such as the Vulcan laser at the Central Laser Facility in the UK[11] and BELLA (BErkeley Lab Laser Accelerator) at the Lawrence Berkeley National Laboratory in the US.[12] These facilities are used to study a wide variety of problems in laser, plasma and nuclear physics and with prospective laser facilities, such as the Extreme Light Infrastructure (ELI), it will be possible to also investigate non-linear QED effects.[13] It should be apparent that these facilities are of great importance for laser-plasma based light ion acceleration and several record-breaking achievements have been made in such facilities.

However, making laser-plasma accelerators a feasible and competitive alternative for light ion acceleration will require more than just sheer power. The efficiency of the energy transfer, from the laser pulse to the light ions, as well as the collimation of the beams will have to be improved. Furthermore, the beams will also have to be made more monoenergetic and the reachable energies will have to be increased further. To accomplish these goals original and clever solutions will have to be devised and perhaps the answer lies in changing the target design, using multiple laser pulses, using alternative field structures of the pulse or in a combination of these. As this can be studied without the record-breaking power output of the largest laser facilities in the world, smaller systems, but still of ultra-high intensity, becomes equally important. The terawatt laser at the Lund High-Power Laser facility in Sweden is one such laser system.[14]

1.2 Problem statement

The aim of this project is first and foremost to get accustomed to the field of ultra intense laser-matter interactions and particularly in performing large-scale simulations in this field, using the Particle-In-Cell (PIC) code PICADOR specifically dedicated to this purpose. In

reaching this goal the project has been divided into two parts, mainly focusing on issues of numerical and physical nature respectively. More specifically, the effects of artificially charging the simulation boundary, when performing Target Normal Sheath Acceleration (TNSA) simulations using PICADOR, are to be investigated and numerical tools to reduce unphysical effects are to be developed if necessary. Furthermore, TNSA simulations are to be performed using alternative target designs. The designs are to be compared on the basis of the resulting light-ion beam characteristics and their performance discussed with regards to key TNSA processes.

1.3 Structure of the thesis

Chapter 2. We begin by providing the reader with a very soft introduction to laser-plasma interactions from a theoretical point of view. In particular, key aspects of target normal sheath acceleration will be discussed.

Chapter 3. In this chapter we will describe the main features of PIC schemes, focusing on their use in simulating laser-plasma interactions. The presented material in this chapter heavily guided by standard literature in the area.[16,17] We will also briefly present the PIC code PICADOR which has been the code of use in this thesis.

Chapter 4. This will be the first of two chapters where we look at what has actually been performed within this project. Here we will look at a couple of numerical issues which are important to either consider or resolve before proceeding with trying to extract physical results from the simulations.

Chapter 5. In this chapter we will turn our attention towards alternative target designs for TNSA, in an attempt to improve the characteristics of the resulting proton beam. The new designs presented here are guided by and chosen such as to confirm previous knowledge of key aspects of TNSA processes, while attempting to keep the designs as simple and realistic as possible.

Chapter 2

Laser-driven plasma-based acceleration

The field of classical electrodynamics has been one of the most successful areas of physics, and science in general, since its birth around the start of the nineteenth century. The theory rests on Maxwell's equations, describing both the evolution and generation of electric and magnetic fields by charges and currents as well as by each other, and the Lorentz force law, describing the force acted upon a charged particle by the electric and magnetic fields.

Maxwell's equations, written in the vector calculus formalism and in their differential form, reads[18]

$$\nabla \cdot \vec{E} = 4\pi\rho, \tag{2.1}$$

$$\nabla \cdot \vec{B} = 0, \tag{2.2}$$

$$\nabla \times \vec{B} = \frac{4\pi}{c}\vec{j} + \frac{1}{c}\frac{\partial \vec{E}}{\partial t}, \tag{2.3}$$

$$\nabla \times \vec{E} = -\frac{1}{c}\frac{\partial \vec{B}}{\partial t}, \tag{2.4}$$

where $\vec{E}$ and $\vec{B}$ are the electric and magnetic fields, $\vec{j}$ is the current density and c is the speed of light. Equations (2.1), (2.3) and (2.4) are named Coulomb's law (or Gauss' law), Ampère's law and Faraday's law respectively. The final equation, (2.2), is typically referred to as *the absence of (free) magnetic charge* (but occasionally also Gauss' law for magnetism). It should be noted that Gaussian (cgs) units are employed throughout this thesis. However, particle energies are for convenience presented in eV.

The Lorentz force law is written

$$\vec{F} = q(\vec{E} + \frac{1}{c}\vec{v} \times \vec{B}), \tag{2.5}$$

where $\vec{F}$ is the force acted on a particle of charge q and mass m and with the fields evaluated at the position of the particle. The equations of motion under the influence of this force are then summarized as

$$\frac{\mathrm{d}\vec{p}}{\mathrm{d}t} = \vec{F}, \tag{2.6}$$

$$\frac{\mathrm{d}\vec{r}}{\mathrm{d}t} = \vec{v}, \tag{2.7}$$

with $\vec{r}$, $\vec{v}$ and $\vec{p}$ being the position, velocity and momentum of the particle. When studying non-relativistic motions of charged particles ($|\vec{v}|/c \ll 1$) the simple velocity-momentum relation,

$$\vec{v} = \frac{\vec{p}}{m}, \tag{2.8}$$

is sufficient. Otherwise the full relativistic relation,

$$\vec{v} = \frac{\vec{p}}{m}\sqrt{1 + (\frac{\vec{p}}{mc})^2}, \tag{2.9}$$

is required to correctly capture the particle dynamics, as is the case for charged particles interacting with intense laser pulses.

2.1 Plasma parameters

Plasma is colloquially referred to as one of the four fundamental states of matter and can loosely be described as an ionized gas. More accurately it consists of unbound, in the sense that they are not atomically bound, electrically charged particles. Most commonly these particles are free electrons together with at least partially ionised ions, but also neutral atoms or molecules can be present. The dynamics of a plasma can therefore be described with Maxwell's equations coupled with the Lorentz force. However, extra care must be taken when neutrals are present in the plasma. A plasma in which this is of importance is referred to as a collisional plasma.

Plasmas are said to be quasi-neutral, meaning that they are neutral only on a global scale. This property arises due to the fact that a separation of charge will set up a field that in turn affects the freely moving charges with a force. This force drives the charges in such a way as to cancel out the field, thereby creating a quasi-neutral state. In reality this is of course a complex and on-going process but what is of importance is the characteristic time scale of this process. Assuming a non-collisional plasma and that the motion of the ions can be neglected, because of their inertia as compared to electrons, it is straightforward to calculate the oscillation frequency of an electron slab that has been displaced from equilibrium. The resulting frequency is called the electron plasma frequency, ω_{pe},

$$\omega_{pe}^2 = \frac{4\pi n_e e^2}{m_e}, \tag{2.10}$$

where ε_0 is the permittivity of vacuum, n_e the unperturbed electron density, e the electron charge and m_e its mass.

As mentioned, one of the fundamental properties of a plasma is the ability to screen out changes in the electric potential. The characteristic length scale at which the potential from a point charge is screened out by the plasma is called the *Debye length*. Again assuming stationary ions the Debye length, λ_D, is given by

$$\lambda_D^2 = \frac{T_e}{4\pi n_e e^2}, \tag{2.11}$$

where T_e is the electron temperature, measured in energy units.

These two parameters define the spatial and temporal scales of the interactions in a simple, non-collisional, plasma and must be resolved to correctly capture the collective behaviour of the plasma in simulations.

2.2 Laser parameters

An important parameter, when considering high intensity laser pulses, is the normalized vector potential, a_0, given by

$$a_0^2 = \frac{2e^2}{\pi m_e^2 c^5}\lambda_0^2 I_0 = 0.731 \frac{\lambda_0^2}{\mu\mathrm{m}^2}\frac{I_0}{10^{18}\,\mathrm{W/cm}^2}, \tag{2.12}$$

where I_0 is the maximal intensity of the laser pulse and λ_0 is its wavelength. When a_0 exceeds unity, the oscillations of an electron in the laser field become relativistic.

For a laser pulse which is Gaussian in both space and time, the laser peak intensity is given by

$$I_0 = \frac{2P}{\pi w_0^2}, \tag{2.13}$$

where w_0 is the waist of the focal spot (the radius at 1/e of the electric field) and with

$$P = 2\sqrt{\frac{\ln 2}{\pi}}\frac{E}{\tau_0}, \tag{2.14}$$

where E is the energy contained in the pulse and τ_0 is the pulse duration defined as the full width at half maximum of the pulse intensity.

Now, if the (electron) plasma frequency, ω_{pe}, is smaller than the laser frequency ω_0 then the characteristic time scale of the plasma is longer than the optical period of the incoming radiation. The medium can not stop the propagation of the electromagnetic wave. The medium is therefore said to be transparent, or under-dense. Concerning the opposite case, when the plasma frequency is greater than the laser frequency, then the characteristic time scale of the electrons is fast enough to adapt to the incoming wave and reflect the radiation at least partially. The medium is said to be over-dense.

These two domains are separated at frequency ω_0, which through the use of Equation (2.10) corresponds to a critical density,

$$n_c = \frac{\omega_0^2 m_e}{4\pi e^2}. \tag{2.15}$$

2.3 Target normal sheath acceleration

One of the today most commonly investigated acceleration schemes for acceleration of light ions is the so-called target normal sheath acceleration scheme, or TNSA for short. In this scheme a very thin solid target is used, typically a foil of micrometer thickness. This foil is then irradiated by an ultra-intense electromagnetic pulse, generated by a high-intensity laser, eventually resulting in the acceleration of light ions. Because of their large mass, very high intensities are required to accelerate the ions directly with the laser pulse. In TNSA, this is circumvented by using electrons in an intermediate step to more efficiently transfer the laser energy to the light ions.

Laser pulses are commonly depicted and treated as having a simple shape, typically Gaussian. In reality this is of course not generally the case and for very intense laser pulses this fact is especially important. A sketch of a more realistic laser pulse can found in Figure 2.1 in which the pulse is shown to consist of two major parts, the main pulse and the prepulse, differing in both intensity and duration by several orders of magnitude. Although the prepulse is substantially less intense than the main pulse, its intensity is still very high and it can therefore interact strongly with its surroundings. The duration of the prepulse is also very much longer than the main pulse, making the interaction time proportionally longer and thereby allowing the prepulse to make substantial changes to

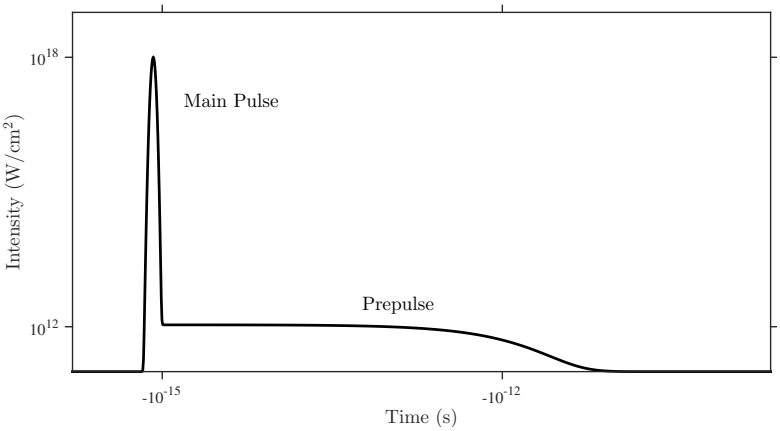

Figure 2.1: *Intensity profile of a (more) realistic pulse consisting of a main pulse and a prepulse, with both an intensity and a duration difference of several orders of magnitude. The time axis is centered at the peak of the main pulse.*

the target before the main pulse reaches it. With the very intense laser pulses used for TNSA, the prepulse ionizes the target which has time to thermally expand before the main pulse arrives. An overdense plasma is formed with a region of so-called preplasma, defined as being the expanded part of the target, where the density gradually increases.

The key features of the TNSA process is depicted in Figure 2.2. After the preplasma has formed and the main pulse reaches the target, highly energetic (hot) electrons are generated at the target surface, around where density reaches the critical density, and pushed into the target at ultra-relativistic speeds. Due to being ultra-relativistic the foil will be transparent to the hot electrons, allowing them to pass through the foil with little energy losses. As they pass through the rear surface of the target a very strong quasi-static electric field will be formed due to the separation of charge at the plasma-vacuum boundary. This charge-separation field will then accelerate positive charges present at the rear surface.

Because of the prevalence of water vapour and hydrocarbon contamination, light ions, ionized by the charge-separation field, will be present at the rear target surface. It is mainly these ions that gets accelerated by the charge separation field, as light ions (and protons especially) typically have larger charge to mass ratios than heavy ions. The acceleration of heavy ions is further suppress as the light ions will then shield the charge-separation field and so a sheath of light ions will be accelerated. Furthermore, this sheath will be accelerated in the direction normal to the rear surface of the target as this is the direction of the charge-separation field, even when the laser is obliquely incident on the target, thus giving rise to the name; target normal sheath acceleration.

Because of the widely differing timescales of the prepulse and main pulse, during which time a multitude of different physical phenomena takes place, it is not feasible to consider the entire chain of events, from ionization to ion acceleration. Instead, TNSA simulations are most commonly initialized just moments before the main pulse reaches the target, which is already ionized and with the preplasma already formed. Furthermore, the light ions meant to be accelerated are typically added as a thin sheath of plasma at the rear side of the target. With the thickness and density of the preplasma and impurity sheaths not easily known or controlled, one typically resorts to educated guesses in determining these parameters for simulation purposes. As for the shape of the preplasma, it is commonly modeled by an exponential or a quadratic function. As an example, the initial density profiles of electrons and light ions are presented for a typical, one-dimensional TNSA simulation in Figure 2.3, with the preplasma here being modeled by a quadratic function. Further details are left to subsequent chapters.

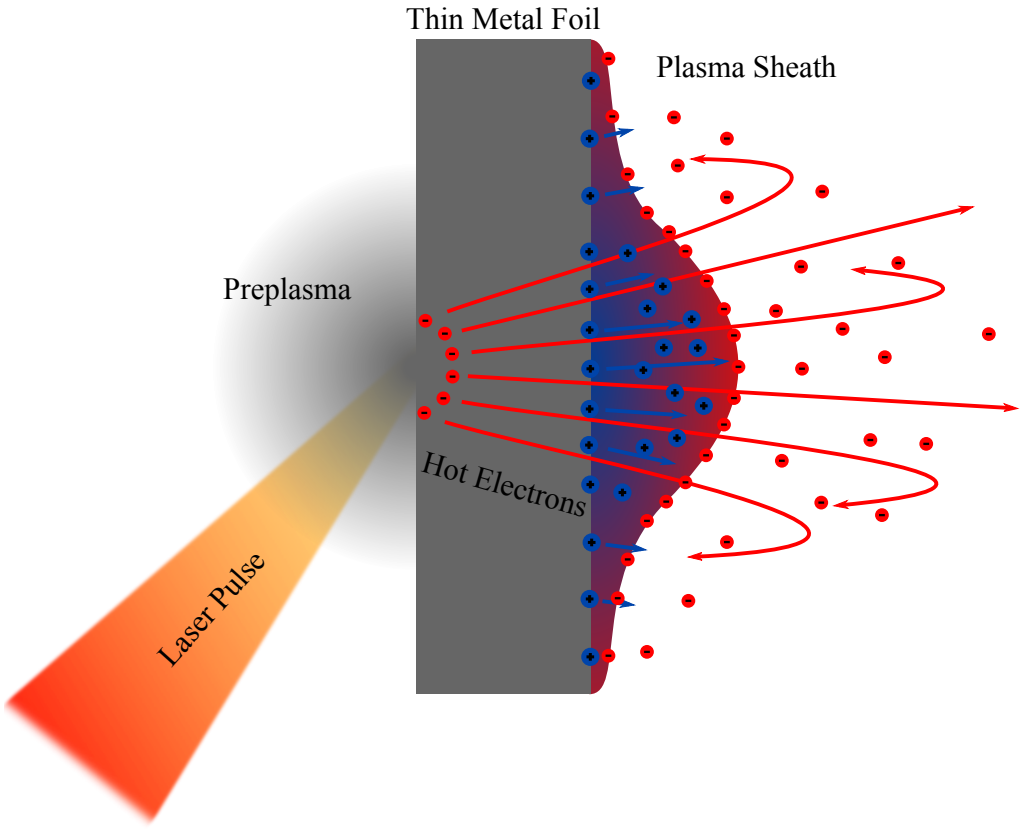

Figure 2.2: *A simplistic illustration of how light ions are accelerated in the TNSA regime using an intense laser pulse.*

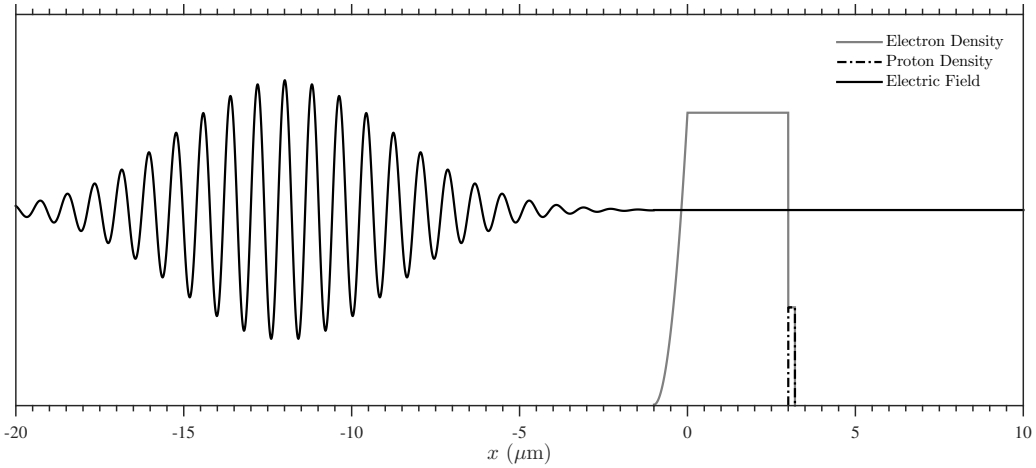

Figure 2.3: *A one-dimensional view of a typical TNSA simulation, with a Gaussian pulse impinging on a target consisting of three charged species: electrons, protons and ions (not shown).*

Chapter 3

Particle-in-cell simulations

As the equations describing the interactions of plasmas and the electromagnetic fields are inherently complex and nonlinear the possibility of studying these systems with analytical tools is quite limited. For relativistic plasmas, such as those induced by ultra-intense laser fields, this possibility is reduced even further. This is especially true when considering problems of more dimension than one, where geometrical considerations can have significant effects on the processes. The need for numerical tools in these areas should therefore be apparent.

Several numerical methods exist for studying plasmas and how they interaction with electromagnetic fields. These techniques are mainly specialized at solving the equations stemming from the different descriptions of plasma physics, each with its own advantages and drawbacks and each requiring different numerical solutions to achieve numerical accuracy and efficiency. The kinetic and fluid descriptions of plasmas are the two general areas on which most numerical simulations are based. In the fluid description the magnetohydrodynamic (MHD) equations of a plasma are solved, assuming approximate transport coefficients. More detailed plasma models, involving particle interactions through the electromagnetic field, are instead considered in the kinetic description.

In this work we only concern ourselves with the particle-in-cell (PIC) method, a technique widely used for kinetic non-collisional plasma simulations. In this method, the electromagnetic fields as well as moments of the distribution such as currents are computed on a stationary mesh with individual particles simultaneously being tracked in a continuous phase space. The strength of this method is that it only samples the phase space, instead of solving the kinetic equations, thereby making it computationally efficient at the cost of some accuracy. Furthermore, this sampling is easily tuned by the number of simulated particles and the PIC method exhibits linear scaling of computational cost with both the number of particles and the number of mesh points, if implemented properly. These properties, together with the fact that the PIC method can easily be parallelized, makes it well suited for large-scale plasma simulations of full dimensionality (6D phase space) on modern supercomputers.

The alternative for studying intense laser-plasma interactions, in which the fluid description is not generally valid, is to solve the plasma kinetic equations, e.g. the Vlasov or Fokker-Planck equations. While the solution of these equations gives a more complete evolution of the phase space and with better precision, as it is not prone to the effects of discrete particle noise and numerical heating present in PIC schemes,[16,19] it becomes impractical for performing many and large simulations of high dimensionality simply because of the high computational cost. Codes in which these equations are solved are therefore mainly used to study problems of reduced dimensionality.

3.1 Basics of the PIC scheme

The implementation of the PIC method is straightforward for plasma simulations, owing to its simplicity in directly simulating the charged particles interacting with and via electromagnetic fields, simply described by the Lorentz force and Maxwell's equations. The scheme can most easily be described by a central loop, in which the simulations are advanced in time. A full cycle in this loop is, for the case of plasma simulation, typically comprised out of four main parts.

- Calculate the currents on the mesh by summing over the contributions of each particle. This is referred to as current deposition and is described in more detail in section 3.1.4.
- Calculate the electric and magnetic fields on the mesh, using Maxwell's equations. This is the so-called field solver, further discussed in section 3.1.3.
- Interpolate the electric and magnetic fields from the mesh to the particle locations. This is referred to as the field interpolation step, further discussed in section 3.1.5.
- Integrate the equations of motion, i.e. the Lorentz force, for each particle. This is called the particle pusher and a more detailed description can be found in section 3.1.2.

The division of the scheme into these fundamental steps not only helps in making the scheme more intuitive, but it provides a good basis for building flexible numerical codes where several different numerical solutions can be provided for each of the steps, the development of which can be made independently. For a schematic diagram of the main loop in a typical PIC scheme, see Figure 3.2.

It should be noted that this section is heavily guided by standard literature on the topic of plasma simulations in general and PIC simulations in particular. The interested reader is therefore advised to consult these works, where a more detailed and complete presentation can be found.[16,17]

3.1.1 Super-particles

As the physical systems typically studied are very large and includes extremely many real particles it is often unfeasible to simulate the system in its entirety. The workaround used in PIC schemes is to have every *simulated* particle be a so-called super-particle, in turn representing a collection of real particles. The number of real particles represented by each super-particle can typically be in the millions for plasma simulations. This method of increasing the efficiency is allowed as it preserves the charge to mass ratio, keeping the Lorentz force unaltered, with the super-particles following trajectories identical to the real particles as a result.

3.1.2 The particle pusher

The particle pusher is, as previously mentioned, the part of the main loop where the equations of motion, given by the Lorentz force, are integrated for all of the simulated particles. There are several, both explicit and implicit, schemes for performing this integration with varying efficiency and accuracy. The most widely used scheme for advancing charged particles is the so-called Boris scheme, an explicit algorithm with excellent long term accuracy. Other common alternatives are the explicit Leapfrog method or the implicit method of Tajima.

The advantage of explicit schemes are that they are generally both simpler and faster than implicit ones, with the typical drawback of less accuracy. However, the particle pusher typically takes a significant portion of the computational time, making it an important step

in attempts at increasing the computational efficiency. The reason for this is that, despite the use of super-particles, the number of simulated particles is usually very large. This is necessary for statistical reasons, as many particles are required to sufficiently sample phase space.

3.1.3 The field solver

The field solver is where the electric and magnetic fields are advanced in time and is entirely done by solving Maxwell's equations on a stationary mesh. Being a system of (hyperbolic) partial differential equations (PDE), the full range of numerical techniques developed to solve this class of PDEs can be utilized. These techniques can be categorized into three main types.

In spectral methods, such as the fast fourier transform (FFT), the system of equations is turned into an eigenvalue problem that is solved globally by using trigonometric basis functions. This method is therefore, contrary to the other methods, not actually solving the equations on a discrete mesh. This is for example used in ELMIS (Extreme Laser-Matter Interaction Simulator), a PIC code entirely based on the fast fourier transform and built for laser-plasma studies.[15]

The second type is finite element methods (FEMs), where the equations again are turned into an eigenvalue problem. As opposed to the spectral methods, this is now done on a discrete mesh of elements generally using low order polynomials as basis functions.

In finite difference methods (FDMs) the discrete mesh is typically called a grid, on which the system of PDEs are turned into algebraic equations by approximating the derivatives as differences between neighbouring grid points. For the solution of Maxwell's equation using finite difference methods the most commonly used technique is the finite difference time domain (FDTD) method. This is the case for the PIC code Picador, to be discussed in further detail in section 3.2.

3.1.4 Current deposition

In this step the current source term is computed on the mesh, as it is needed for solve Maxwell's equations correctly in the field solver and without the current calculated in this step the particles would not generate any fields in the simulations. The deposition is performed by assigning a shape, or *form factor*, to each simulated particle. This form factor then determines how the current, as generated by the super-particle, is distributed over the adjacent grid points. Commonly used form factors are the so-called nearest grid point (NGP), cloud-in-cell (CIC) and triangular shaped cloud (TSC). These are zeroth, first and second order weighting schemes respectively.

Furthermore, there also exists specialized schemes such as Esirkepov, which can be implemented with different orders of convergence as well. What is special about this scheme is that it takes care of the problem of charge conservation, which is not assured with the simpler schemes.

3.1.5 Field interpolation

To integrate the equations of motion, for each particle, in the particle pusher the electric and magnetic fields have to be known at the position of each particle. However, the fields are calculated on a discrete mesh in the field solver, whereas the particles are located in continuous space. Thus, the fields at the position of each particle must be computed from the fields on the mesh. This is the purpose of the field interpolation step in the loop.

As in the current deposition schemes, the field interpolation is done using form factors. This has to be performed such that the fields at different grid points contribute to the fields at the particles in a manner consistent with Maxwell's equations. The weighting

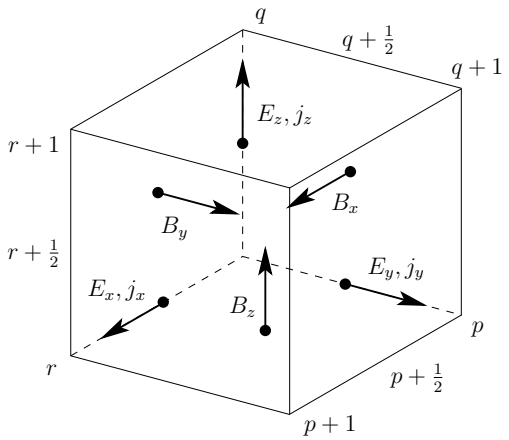

Figure 3.1: *A representation of the Yee grid unit cell. The indices p, q and r indicate node positions in the x, y and z directions respectively.*

scheme is typically the same as the one used for the current deposition, as the PIC scheme can be made to conserve momentum this way.

3.2 PICADOR

The particle-in-cell scheme is a central tool to the work done in this project and it is through the numerical code PICADOR that this is being used. PICADOR is a PIC code written in C++ and designed specifically for the purpose of intense laser-matter studies.[20] It is fully parallelized using both MPI and OpenMP and is therefore well suited for large-scale plasma simulations on modern super-computing clusters. This need for high-performance computing comes from the fact that contemporary problems typically require simulation of $\sim 10^8 - 10^9$ particles and meshes with $\sim 10^7 - 10^8$ grid points and above.

PICADOR is developed by a collaboration of researchers from Lobachevsky State University of Nizhni Novgorod, the Institute of Applied Physics of the Russian Academy of Sciences, Chalmers University of Technology, Umeå University and Lund Laser Center.[21] The parallelization of the code is done by dividing the simulation region into spatially uniform subdomains, but allowing for different divisions in different spatial dimensions. Each of these subdomains will have an integer number of grid cells that makes up the discrete mesh. The subdomains are then handled by one MPI process each, while threads work on individual cells in cases where OpenMP is used. Apart from being fully parallelized, PICADOR also supports dynamic load balancing,[22] in which the relative workload of different MPI processes is dynamically changed throughout the simulation by means of rectilinear re-partitioning of the simulation subdomains. As the rectilinear decomposition is topologically equivalent to the spatially uniform decomposition the interactions between MPI processes remain unchanged even though different processes may handle differently sized subdomains, thus giving the opportunity of increasing the computational efficiency for problems with non-uniform particle distributions.

Simulations in PICADOR are centered around the main loop of the PIC scheme and features FDTD and NDF field solvers, two finite-difference solvers that solves Maxwell's equations on a staggered grid, called a Yee grid. This grid uniformly covers the simulation region and the Yee grid unit cell can be seen in Figure 3.1, where it should be noted that the electric and magnetic fields are computed not only on displaced grid points but at different instants in time (differing by half a time step). The current density is computed in nodes at the edges of the grid cells, as is the electric field, whereas the magnetic field is

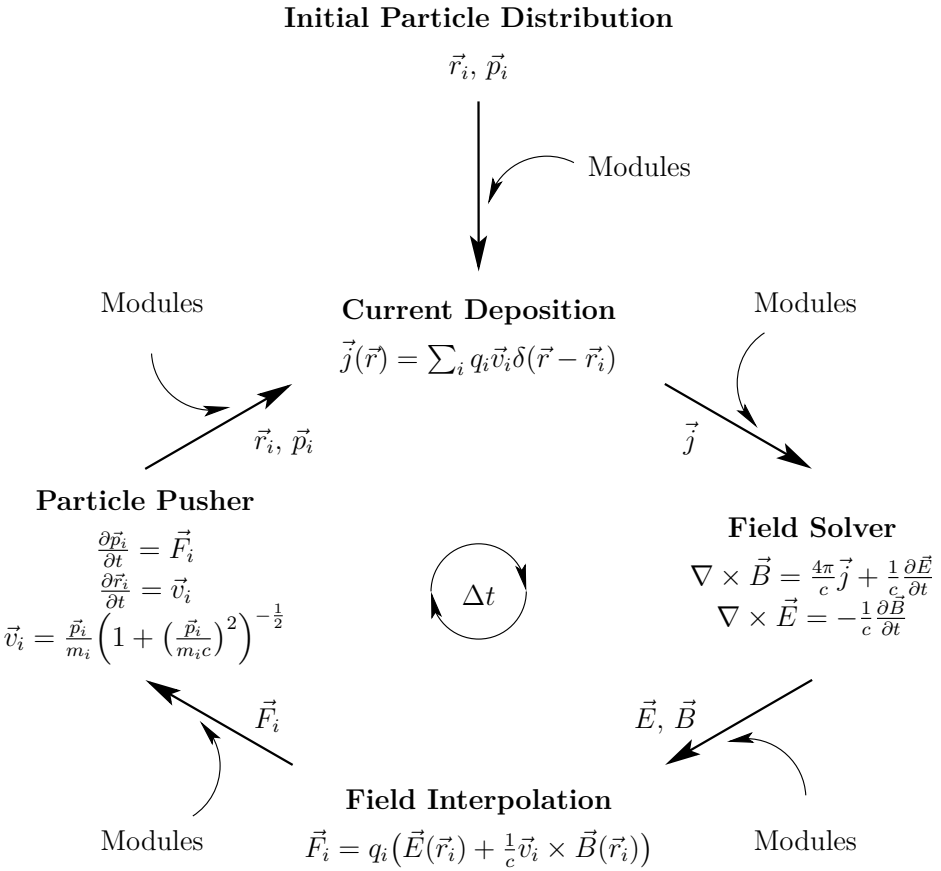

Figure 3.2: *A schematic diagram of the main loop in a typical particle-in-cell code, for the purpose of plasma simulations.* PICADOR *allows flexible extension of the PIC scheme through the use of modules that can be included at different points in the loop and can be independently activated for separate simulations.*

computed on nodes located at the grid cell faces.

In computing the fields with the FDTD and NDF solvers only Ampère's and Faraday's laws, Equations (2.3) and (2.4), are needed. The reason for this is a neat property of the Yee grid, namely that it identically satisfies Gauss law for magnetism, Equation (2.2). Furthermore, if the PIC scheme is charge-conserving, as is the case in PICADOR through the use of Esirkepov current deposition, the last of Maxwell's equations, Equation (2.1), will merely be an initial condition for the simulation. As simulations typically considers quasineutral plasmas, where electrostatic fields are not initially present, Poisson's equation is most often already satisfied (with initial fields being zero as default) and therefore not considered. However, it is possible to manually set the initial fields in cases where this do not hold.

The main loop further features the Boris particle pusher, CIC and TSC particle form factors for field interpolation and current deposition as well as Esirkepov current deposition. Apart from this, PICADOR is built with the possibility of including extended features through so-called *modules*, as depicted in Figure 3.2. Dynamic load balancing mentioned previously is one such feature and others include ionization, moving frame, QED event generation and absorbing boundary layers. Furthermore, PICADOR is fully equipped to handle both periodic and absorbing boundary conditions, independently chosen for different spatial dimensions.

Chapter 4

Numerical issues

This is the first of two chapters where the work done in this thesis is presented. Here we focus our attention on numerical aspects of laser-plasma PIC simulations by dealing with a number of numerical issues in PICADOR. We will here make use of the possibility of adding extended features to PICADOR by developing new modules for diagnostics, in an attempt at identifying the cause of specific issues, as well as for resolving identified issues or suppressing found unphysical effects. Exact details of the developed modules are however not presented in this thesis.

4.1 Choosing numerical parameters (basics)

When performing PIC simulations there are a number of important parameters that one needs to know and specify from the beginning of the simulation, the most obvious parameters being the time and space steps Δt, Δx, Δy and Δz. As fixed spatial and temporal scale are required to cover the main aspects of the physical system to be simulation, the choice of these parameters will of course affect the computational cost of performing the simulation as well as the precision and accuracy of the simulation. These parameters can not, however, be taken to be too large as there are limitations both due to physical and numerical considerations, wherein the resolution of certain physical effects must be sufficient and the numerical schemes must be kept stable.

Other parameters of importance are related to the initial conditions of the physical system to be simulated. For instance, the initial temperature and density of the plasma. These parameters obviously change the behaviour of the plasma which in turn can affect the stability of the simulation. As an example, increasing the electron plasma density will affect both the electron plasma frequency, Equation (2.10), and the Debye length, Equation (2.11), potentially requiring an alteration of the time and space steps, such that the key features of the plasma remain resolved.

When studying the spatial resolution of the simulation, there are mainly two things that are important to resolve. The first thing is the wavelength of both the initial and in the simulation generated radiation. The second is the already mentioned Debye length. For all intents and purposes, we will in this work always be located in parameter regimes where the Debye length is the smallest length scale of the system.

Turning to the temporal resolution there are, again, mainly two things necessary to resolve in the simulation. The first being the plasma period, given by the plasma frequency, and the second being the frequency of both the initial and in the simulation generated radiation. Assuming that the resolution of the wavelength of the radiation has already been dealt with the resolution of the frequency now becomes trivial, seeing that the propagation

speed is fixed with the speed of light.

$$\Delta t \lesssim \frac{\Delta x}{c}, \quad \text{assuming } \Delta x = \Delta y = \Delta z. \tag{4.1}$$

However, there is actually another thing that needs to be kept in mind when dealing with the temporal resolution and which will make the estimation in Equation (4.1) obsolete. For the FDTD scheme to be stable and converge it has to fulfill the Courant–Friedrichs–Lewy (CFL) condition.[23] The CFL condition is a necessary, but not necessarily sufficient, condition for convergence when numerically solving (hyperbolic) partial differential equations using the method of finite differences. For the case of the FDTD scheme, the CFL condition is

$$\Delta t \leq \frac{1}{c\sqrt{\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2}}}, \tag{4.2}$$

which for the case of using cubic unit cells ($\Delta x = \Delta y = \Delta z$) becomes

$$\Delta t \leq \frac{\Delta x}{c\sqrt{3}}. \tag{4.3}$$

The question then becomes which of these two, the CFL condition or the plasma period, defines the smallest time scale? Again, for this particular work we will be working with parameters where the CFL condition defines the smallest time scale in the system.

Furthermore, before any simulation can be performed the size of the simulation area has to be decided, by choosing the number of grid cells in each dimension, and in case absorbing boundary conditions are required, the size of the field-absorbing boundary layer has to be chosen. The size of this layer is usually taken to be on the order of the wavelength of the radiation to be absorbed such as to sufficiently absorb the fields, but without taking up too much of the simulation area, with the laser wavelength typically being a good reference for the purpose of laser-plasma simulations.

4.2 Measuring numerical effects

After it has been determined what defines the spatial and temporal scales in the simulation comes the tricky questions. What does it actually mean to resolve the Debye length? It should be obvious that simply taking the space steps to be at least one order of magnitude smaller than the Debye length should be sufficient. However, decreasing the space step would also require a larger grid to encompass the same simulation region, thus increasing the computational cost. For this reason we will need some way of deciding when the space step is sufficiently resolved.

As it turns out, studying the total energy in the simulation region, both the kinetic energy of the super-particles and the energy in the electromagnetic fields, is a good way of monitoring unphysical effects. In general, the FDTD scheme by itself conserves the energy in the fields given that the CFL condition is met, in the opposite case the scheme is unstable and the energy increases exponentially. However, this is not generally true when coupling the field solver to a particle mover and PIC schemes often experience numerical heating effects.[16] These effects can be decreased by increasing the resolution, as this in turn will decrease the particle factor. The particle factor is important in this regard since the heating and cooling effects comes from so-called discrete particle noise. This can then of course also be alleviated by decreasing the plasma density as, all other things equal, this will reduce the particle factor as well. However, this would then of course also change the physics of the simulated system and may therefore not be a feasible alternative.

This energy analysis becomes even more difficult when considering systems where energy can both enter and leave the simulation region via for example absorbing boundaries for

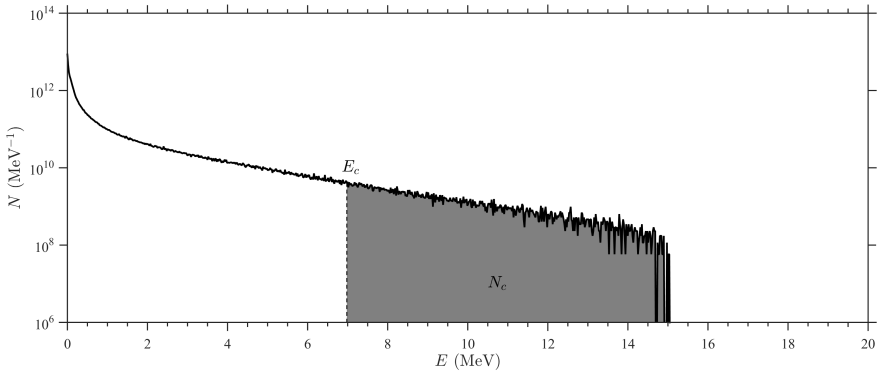

Figure 4.1: *Cutoff energy, E_c, and particle number, N_c for a sample spectrum. The cutoff energy is defined such that N_c particles have an energy above, or equal to, this cutoff.*

both particles and fields. A more complex solution would in these cases have to be considered, if an accurate analysis based on the system energy is to be made. A proposal for the further development of this energy diagnostics is therefore presented in section 4.2.5.

It can of course be argued that the best way of ensuring the accuracy and validity of the numerical simulations is to study the numerical schemes more rigorously by looking at for example convergence towards exact solutions. This should of course be done as well, but when the numerical scheme gets very large and complex, as is the case for PICADOR with its many module extensions, it becomes difficult to cover all possibilities. A simpler method is then to monitor well-understood physical quantities to see if they behave as expected. The total energy of a physical system is one such quantity with others being the total momentum and the total charge, quantities that are typically conserved in closed systems. This is of course not a guarantee that a simulation will provide the correct solution but can be used as a first step of controlling simulations, and with little effort.

Besides monitoring of the total energy we will employ a different method, aimed more specifically towards TNSA simulations. What we want is some simplified way of describing the properties of the resulting proton beam. More specifically, we want describe the resulting energy spectrum with only one, or a few, numbers. This can of course be done in multiple different ways, all with different benefits and drawbacks, but the method employed here was chosen because it besides being simple, quite general and could be of benefit when discussing ion beam requirements for different applications. The way this will be done is by using the definition

$$N_c = \int_{E_c}^{\infty} \frac{\partial N}{\partial E} \, \mathrm{d}E, \tag{4.4}$$

where N_c is the cutoff particle number, E_c the cutoff energy and $\frac{\partial N}{\partial E}$ is the particle energy spectrum. This is also presented in Figure 4.1 using an example spectrum. The definition can be used in two different ways, either by deciding upon a cutoff energy and calculating the corresponding particle number or by deciding upon a particle number and calculating the corresponding cutoff energy. In this work, the latter will be most extensively employed.

As with all conceivable methods of describing an entire spectrum with only one number, this method has the drawback of not being able to distinguish between identical results arising from spectra of different shape, although some might be more desirable. As such, a fair comparison between different spectra requires that the spectra in question are of similar shape, as is the case for the proton energy spectra considered in this work. In the event that the spectra are very different the best method would most likely be to simply show the entire spectra.

4.2.1 Numerical heating and instability

Upon performing two dimensional simulations using PICADOR it was found that the total energy in the simulation volume increased considerably. Further studies, presented here, indicated that this could not be explained by means of ordinary numerical heating, but that it must be some numerical instability. It was possible to mitigate this increase by a trade-off in computational costs, but only the extent of making reasonable – physics wise – two dimensional simulations unfeasible.

In an attempt at getting to the bottom of this apparent heating issue a numerical study was performed, focusing on the energy change per unit volume in a uniform, infinite, plasma and how different choices of simulation parameters affected this rate of heating. In Figure 4.2 you will find the relative change, within 100 cycles of the plasma period, of the total energy in the simulation volume presented against the resolution of the Debye length and for a number of different parameter choices.

At this stage it is only important to discuss what can be seen from all but the last parameter choice (green crosses), which we will return to in due time. Upon increasing the initial plasma temperature, for both choices of temporal resolution, we observe that the energy change also increase. The exact reason for this remains unclear. However, there is no real reason to suspect that making this change should have no effect on the ordinary numerical heating and as it in no other way changes the rate of energy increase, it is reasonable to suspect that the cause of our heating issue lies elsewhere.

Instead looking at how a decrease in temporal resolution, for both choices of initial plasma temperature, affects the rate of energy increase we find that a drastic change occur for $\Delta x < \lambda_D$. There are two main things one can identify as problematic with this. First and foremost, there is no real reason to suspect that such a drastic change should occur by simply reducing the temporal resolution, if there were no instabilities involved. Secondly, both parameter sets with the decreased temporal resolution initially show a decreased rate of heating for an increased spatial resolution. However, this trend eventually stops and turns into a trend of rising rate of heating, for an increasing spatial resolution. It could be argued, against the first of these points, that the CFL condition is only a necessary requirement and that the chosen temporal resolution is simply too close to the limit. While this could indeed be true, the instability due to not reaching a sufficient convergence in the FDTD scheme is typically more extreme than what is observed here, blowing up quickly in only a small number of time steps.

Further on it was found that this issue was also dependent upon the choice of space step in the z-direction Δz, even though the simulations were supposed to be two-dimensional. It was also found that the CIC and TSC current deposition schemes were highly unstable, as the energy increased by extreme amounts when using these schemes instead of Esirkepov. All things combined, the most likely explanation was therefore that the issue came from some numerical instability, rather than the ordinary numerical heating effects present in PIC simulations. It also indicated that this instability could not be traced back to trivial reason such as simply bad choice of spatial and temporal resolution. One possibility being that the instability was rather due to incorrect implementation somewhere in the code.

With what was now known about the issue two separate implementation errors were identified. The first one was caused by the fact that the current deposition for two-dimensional simulations was done in the same way as for for three-dimensional ones. This is possible to do in terms of the grid, as extra so-called ghost cells are used. However, as a result, the current created by each particle was spread along the z-direction according to the particle form factor when it, for 2D simulations, should actually be independent of the z-coordinate. This in turn has the effect of making the fields z-dependent as well, but as periodic boundary conditions are used in the z-direction this eventually leads to our heating issue. To remedy this the currents are instead computed via the sum along the z-direction and then assigned to all cells, for each x, y position. Furthermore, the

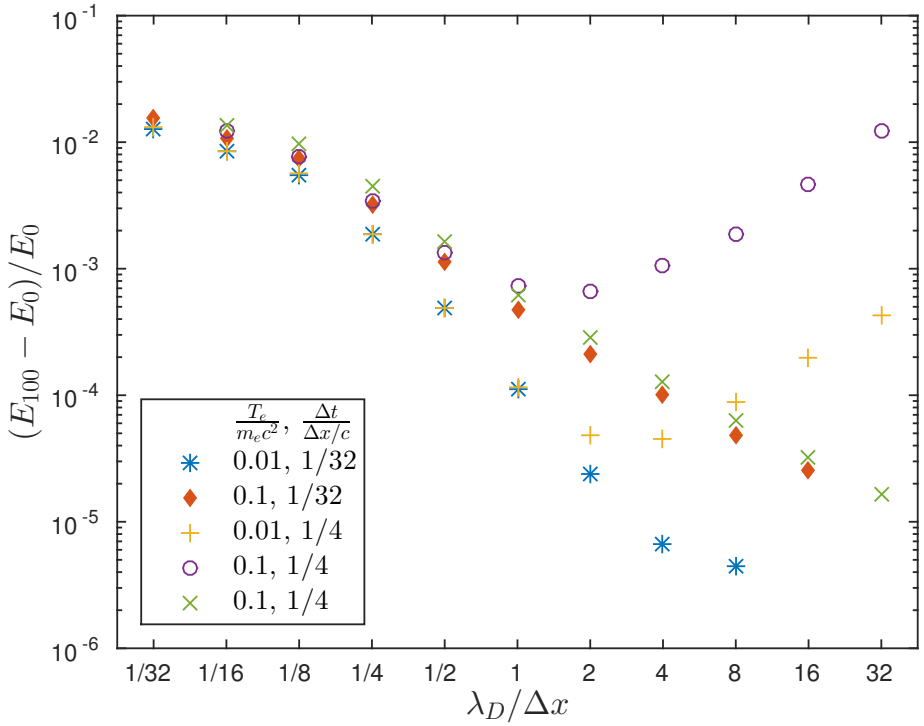

Figure 4.2: *The figure shows the relative change of the total energy in the simulation volume, in 100 cycles of the plasma period, for different resolutions of the Debye length. This is done for two different initial temperatures and two different temporal resolutions. The final set of simulations (green crosses) is identical to the penultimate set (purple circles) in initial parameters. However, it shows the same trend as the corresponding set with highly increased temporal resolution (orange diamonds), as the final simulation was performed after the resolution of the heating issue.*

field boundary conditions are made to ensure that the field values are identical in the z-direction, for each x, y position.

The second error is related to the moments in time in which particle positions and momenta are defined. In the combination of numerical schemes used in PICADOR the positions are defined at integer time steps, $i\Delta t$, whereas the momenta are defined at half-integer time steps, $(i+1)\Delta t$. This fact was correctly taken care of in the Esirkepov current deposition scheme, but not in the CIC and TSC schemes, thus making them highly unstable. For these two schemes, the currents are instead computed after the particles are pushed the first half of a time step, so that the particle positions are also defined at $(i+1)\Delta t$. The remainder of the push is then performed and the issue is resolved.

After these changes had been carried out a final simulation set was performed, also presented in Figure 4.2 (green crosses). It can be seen that this simulation set rather than following the, in terms of parameter choices, identical simulation set (purple circles) it follows the trend of the simulation set with a greatly increased temporal resolution (orange diamonds). This is what we would expect to see if the numerical instability was dealt with and it now shows that the temporal resolution has no noteworthy effect on the rate of heating.

To further show that the instability issue has been resolved it is instructive to look at the time evolution of the total energy in the simulation region, for a few different parameter

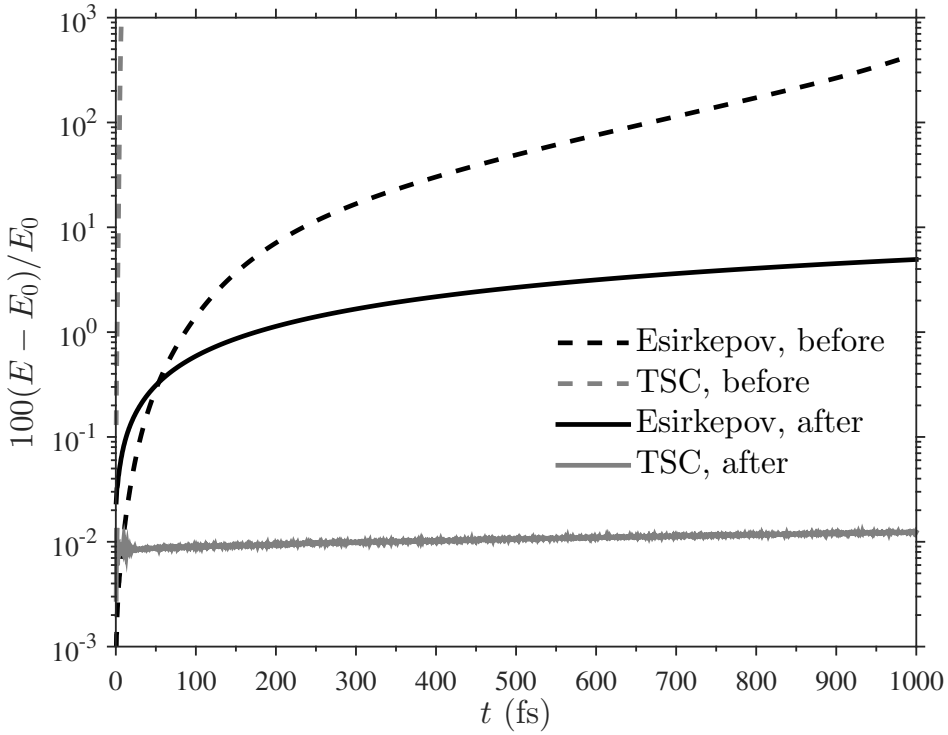

Figure 4.3: *Time evolution of the total energy in the simulation volume for a given parameter sets, using Esirkepov and TSC current deposition both before and after the stability issue has been resolved.*

choices, presented in Figure 4.3. Here it becomes quite apparent both that the simulations before the fix could not reasonably describe the physics accurately and that the simulations after the fix has become stable to a pleasing degree.

4.2.2 Reaching saturation in ion acceleration

When performing simulations of the TNSA process it is the properties of the resulting beam of light ions, or in our case protons, that are of greatest interest. One such property, and perhaps the one of greatest importance, is the kinetic energy of the protons in the beam. To be able to extract this energy, or rather the entire energy spectrum, the simulation will have to be performed over a long enough period of time and the simulation region will, besides this, have to be sufficiently large. This is necessary both to allow the protons get accelerated by the charge-separation field for as long as the field remains and to make sure that the field and particles of the expanding plasma sheath do not get absorbed at the simulation boundary before the acceleration of the protons has stopped.

When the protons have stopped accelerating they are said to have reached *saturation*, after which point it is no longer necessary to keep the simulation going. However, it is not trivial to determine how long the simulation will have to last and how large the simulation box will have to be to encapsulate the entirety of the proton acceleration. While reasonable estimates can be made it depends highly on the specific simulation setup, for example the target geometry, if the simulation time and box size are large enough.

4.2.3 Removing charge from the boundary

The electromagnetic fields generated by the particles are, as is described in chapter 3, accounted for via the current that the particles produce. When a particle reaches and crosses the simulation boundary its contribution to the current is removed. The fields generated due to this particle will therefore behave as if the particle has stopped moving as it crossed the boundary, giving the boundary an artificial charge. This will in turn make fields and particles interact with the boundary in a way that is unintended, potentially yielding incorrect results.

As the acceleration of the protons potentially goes on for the entire distance between the target and the rear simulation boundary, the protons can get strongly affected by any boundary effects as they approach the boundary. It is also at this point in the simulation that we will want to extract the resulting energy spectrum, so any numerical effects at this stage could significantly alter the final result. As a result it was suggested that it should be investigated both whether or not this is of an issue in TNSA simulations with the PIC implementation of PICADOR and if this effect can be removed by simply pushing particles away from the rear boundary, in the transverse directions, before they get absorbed. That is, the particles are given a large momentum in a direction transverse to the rear simulation boundary, and target normal, when they get too close to the rear boundary. In this case, too close means just before entering the absorbing layer.

The setup of a two-dimensional TNSA simulation with a plain target is presented in Figure 4.4. Using density profiles similar to what is presented in Figure 2.3 a number of such simulations were performed and with varying simulation box size in the x-direction, L_x. The resulting cutoff energies are presented in Figure 4.5 for both the ordinary case and with charges being pushed towards the transverse boundaries. The incoming laser pulse had a Gaussian shape with a duration of 40 fs, an energy of 0.75 J and the pulse hit the target at normal incidence. In doing the simulations, the simulation time was taken to be sufficiently large for the expanding sheath to get to the rear boundary and the simulation box size in the y-direction, L_y, was chosen such that effects from the transverse boundaries could be neglected (it is taken to be several times greater than the pulse diameter). From the figure we find that by increasing L_x the resulting cutoff energy eventually saturates. This is well in line with what we expected to see and it also shows, for this particular setup, that a simulation box size above at least 30 µm is sufficient in the x-direction.

Furthermore, the figure also shows the effect of pushing particles towards the transverse boundaries. More specifically, it is only the electrons that are being pushed. As the simulation box size is increased the difference between the two simulation sets decrease and the simulations where the electrons are pushed away invariably gives the lowest cutoff energy. The reason for this should be quite apparent. As the simulation box is increased, the amount of electrons that gets removed at the boundary is decreased as they instead have the opportunity to recirculate back to the target and those electrons that do get removed at the rear boundary should also do so in a less localized region. The result of simply increasing the box size should therefore be that the numerical effects due to the artificially charged boundary should simply decrease. This is well supported by the fact that no significant difference between the two simulation sets can be seen for a large enough simulation box.

4.2.4 Computational costs

The greatest restriction on the accuracy and precision that can be achieved in simulations most often comes in the form of available computational resources. Any changes to the software in use that yields a computational speed-up without affecting or, even better, increases the accuracy and precision of the results, is therefore of great value as this in turn opens up new possibilities as to what can be studied.

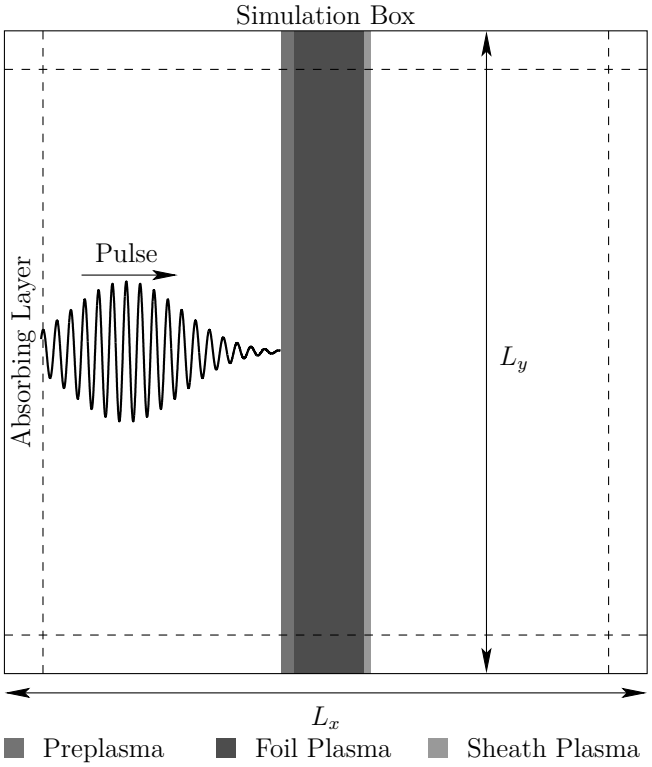

Figure 4.4: *A 2-dimensional view of a typical TNSA simulation, with a Gaussian pulse impinging on a target consisting of three parts: the pre-plasma, the foil plasma, and the sheath plasma. The outermost parts of the simulation region consists of a layer that absorbs the field, to reduce unwanted reflections.*

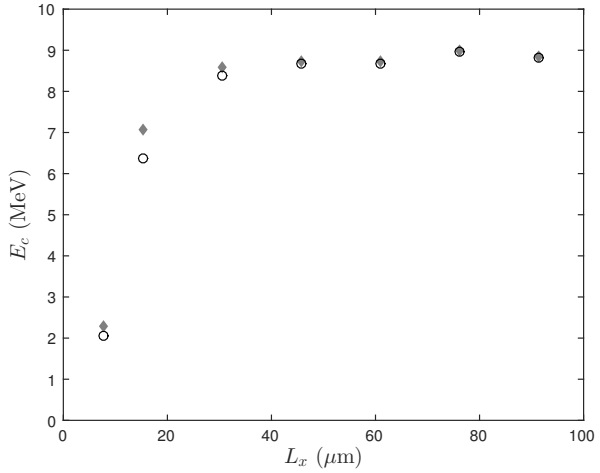

Figure 4.5: *Saturation of the cutoff energy, E_c, using a particle number of $N_c = 10^9$ for a simulation box size of $L_y = 60\,\mu\text{m}$ and varying L_x between simulation. This was done for two different cases: standard TNSA (diamonds) and standard TNSA while removing electrons before the rear absorbing layer (circles).*

While discussing the computational costs of running a reasonably sized simulation it might be interesting to go back to the instability issue and note that the changes made where not only aimed at removing the instability, but also to speed up 2D simulations in PICADOR. How drastic of a change this actually was is most easily described by numbers. We again turn to Figure 4.3 where the time evolution of the total energy is presented for a couple of simulations. The figure shows how the simulations are made much more reliable by removing the unstable behaviour, but what it does not show is that the Esirkepov simulation, before the fix, required 12000 core hours to finish, whereas the TSC simulation, after the fix, only required 1500 core hours. The reason for this is the big difference in resolution that was necessary to reduce the heating. At the same time, the heating was reduced from being on the order of 10^2 percent down to 10^{-2} percent, thereby making it possible to perform reliable 2D simulations in PICADOR on a daily basis.

4.2.5 Suggestions

The idea of using the total energy of the simulation area, as was introduced in section 4.2, has been extensively used throughout this chapter. However, as was noted previously, this diagnostics tool can be made more useful to situations where it may be of importance to also monitor the flow of energy through the boundaries. This can be done by adding the possibility of computing the energy flow through the various faces of the simulation box. This idea rests on Poynting's theorem, which is a statement of the energy conservation for the electromagnetic field,

$$\frac{\partial u}{\partial t} = -\nabla \cdot \vec{S} - \vec{j} \cdot \vec{E}, \tag{4.5}$$

where u is the electromagnetic energy density, $\vec{S}$ the Poynting vector, $\vec{j}$ the current density and $\vec{E}$ the electric field.

The electromagnetic energy density, in the first term of Poynting's theorem, is already used to calculate the total electromagnetic energy inside the simulation box through the relation

$$u = \frac{1}{8\pi}\left(\vec{E}^2 + \vec{B}^2\right), \tag{4.6}$$

where $\vec{B}$ is the magnetic field. The Poynting vector, in the second term of the theorem, is given by

$$\vec{S} = \frac{c}{4\pi}\vec{E} \times \vec{B}, \tag{4.7}$$

and represents the directional energy flux density of the electromagnetic field. The proposal is thus to calculate the Poynting vector at the simulation box faces and in doing so the rate of energy change of the simulation box due to radiation, can be monitored. The potential difficulty in doing this lies in the fact that the electric and magnetic fields need to be computed at the same position, making interpolation a requirement.

The last term of Poynting's theorem can be taken care of in several ways. It corresponds to the rate at which the electromagnetic fields do work on charged particles and is thereby taken care of as the total kinetic energy of the particles is already monitored as well. It would however be straightforward to also calculate this term using the currents and electric fields on the grid, as they are located at the same positions on the grid.

Chapter 5

Alternative targets for improved TNSA

In this chapter we will investigate some very simple ways of improving the characteristics of the light ion beam resulting from TNSA, by using alternative target designs. From a simulations point of view there are no real restrictions on how the target can be designed, it is perhaps only limited by human imagination. However, this of course does not make any setup practical, or feasible. We therefore restrict ourselves to the most basic designs, which we today already know are realizable.

Presented in Figure 5.1 is a schematic diagram over different energy routes and how the energy goes into less favourable channels during a standard TNSA process. It also shows potential ways of reducing the energy lost to these channels, or how to avoid them altogether. Using the ideas presented in this figure, we attempt to increase the energy that eventually finds its way to the high-energy collimated light ions, protons more specifically, at the back of the target. This is done by making the target designs so that they alter at least, but mainly, one of the main energy transfers shown in the figure.

A structured target will be used to decrease the amount of reflected radiation, from the pulse impinging on the front side of the target. To decrease the transport of electrons in the transverse – to the target normal – direction, a target with finite extension in the transverse direction will be used. Finally, an attempt is made to increase the charge-separation field on the rear side of the target, by using the same structured target as before, but rotated 180 degrees such that the structures are now located on the rear side of the target. As mentioned, more exotic targets can of course be analysed using the very same tools and for example the use of a spherical target has already been investigated.[24] For more examples see Macchi et al.[25] and references therein.

The simulations have been carried out using laser parameters of the terawatt laser at the Lund High-Power Laser facility in Sweden.[14,26] More specifically a 40 fs long, linearly polarised, Gaussian pulse with an energy of 0.75 J and a focal spot diameter of 7 μm is used. Furthermore, the pulse has a central wavelength of 810 nm and the intensity peak of the pulse is set to reach the target after 100 fs. The prepulse is treated as described in chapter 3, with a preplasma layer assumed to be already generated with a fixed density profile at the onset of the simulation. The simulations are performed in two spatial dimensions to reduce the computational cost, thus assuming that both the fields and densities extend infinitely and uniformly into the third spatial dimension. Since the simulations are performed in only two spatial dimensions, some sort of trick is required in order to get an estimate on the actual particle numbers a setup can produce. For this purpose, while the simulation assumes a truly two dimensional setup, we will assume that the TNSA process only occurs within a finite range in the third spatial dimension and that this range is on the order of

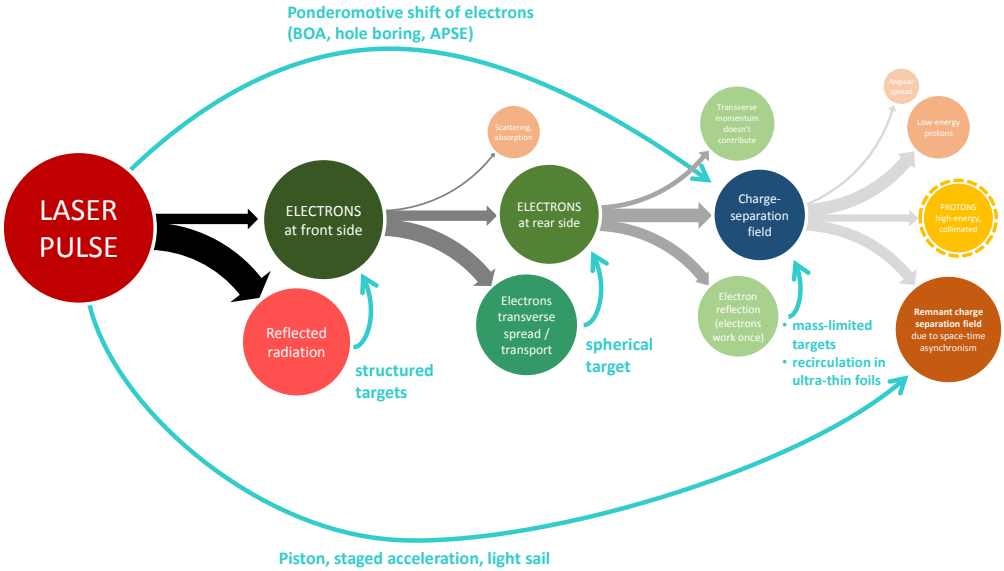

Figure 5.1: *A schematic diagram over different energy routes and how the energy goes into less favourable channels during a standard TNSA process. The blue arrows show potential ways of reducing the energy lost to these channels, or how to avoid them altogether.*

the focal spot size of the pulse. For convenience we take the range to be exactly 10 µm.

As the main objective is to analyse the alternative target setups according to their capabilities in improving the quality of the proton beam, the setups will have to be compared to an equivalent standard TNSA setup. In Figure 5.3 the electric field and the electron and proton densities are presented for a selection of different times into the simulation. The color scales for these figures, and for that of the corresponding figures of subsequent setups, are presented in Figure 5.2. From the figures we see that a strong charge-separation field is beginning to form by the time the peak of the pulse reaches the target, Figure 5.3b. This field continues to grow for the first couple of femtoseconds preceding this frame and by the time of Figure 5.3c, when all but a small fraction of the pulse has been either absorbed or reflected by the target, the charge-separation field has reached its full strength. The strongest part of this field can be seen to travel in the positive transverse direction. This is the effect of the transport of hot electrons in the transverse direction, also visible in the figures, and while a strong field does remain to accelerate the now expanding proton sheath, Figures 5.3d–5.3f, a substantial part of the energy is evidently lost this way.

This anisotropy in the momenta of the generated hot electrons is much due to the homogeneity of the plain target. Consider the case of an intense pulse impinging on a plain target at an angle, θ. It can be found that the momentum components of the generated hot electrons in the x- and y-directions (p_x and p_y) are related by[27]

$$p_x = \frac{\sin\theta}{\cos^2\theta}\left(\sqrt{1+p_y^2\cos^2\theta}-1\right), \tag{5.1}$$

from which it can be seen that the 'hottest' electrons will travel in the same direction as the initial direction of the pulse. The calculation builds on the fact that the target is homogeneous in the y-direction and the anisotropy can thus be altered by breaking this homogeneity of the target, thereby reducing the unwanted transport of hot electrons in the transverse direction.

The evolution of the proton energy spectrum is presented in Figure 5.5a along with

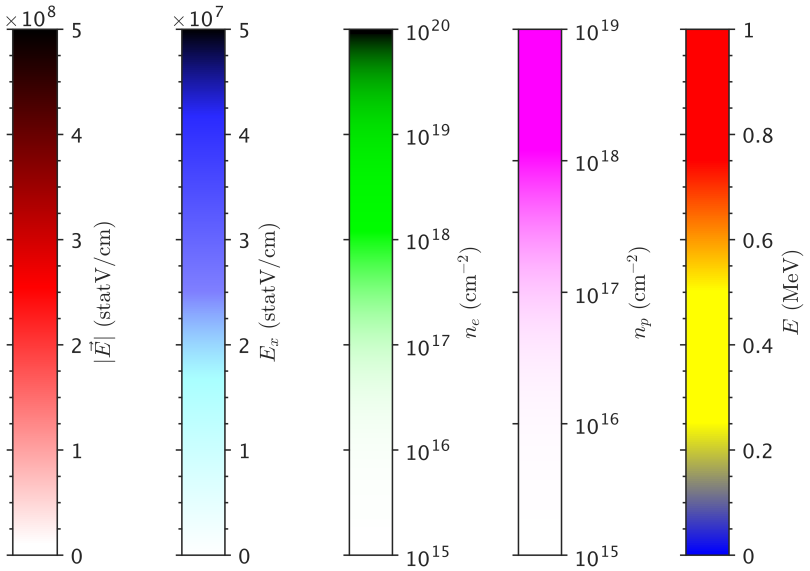

Figure 5.2: *The color scales of, from the left, the absolute value of the electric field, the normal (to the target) component of the electric field, the density of electrons, the density of protons and the energy of tracked electrons are here presented for Figures 5.3, 5.5, 5.7, 5.9 and 5.11.*

the time evolution of different energy cutoffs. From this figure we see that the spectrum over time gets saturated and that the charge-separation field act such as to give us a very broad spectrum by only being able to affect the protons of highest energy, present at the very front of the expanding sheath. The rapid change in the spectrum after and around the 1000 fs mark is due to the fact that the most highly energetic protons start leaving the simulation region and the simulation can therefore not be considered accurate any longer.

Also of interest, when discussing the proton beam quality, is the angular distribution of protons at different energies, as is presented in Figure 5.5b for a given instant in time. From this figure it can be seen that the protons are traveling mainly in the normal direction, only deviating slightly from the normal by a few degrees and with a spreading on the same order.

As is easy to imagine, the two last figures just discussed will be especially important when comparing the alternative target designs to this plain, standard target. We will therefore return to these on numerous occasions in the succeeding sections.

Figure 5.3: *Simulation of a standard TNSA setup, with an pulse of 0.75 J incoming at an angle of 45 degrees to the target normal. The target consists of a 3 µm thick foil of plasma initially described by a density profile as shown in Figure 2.3. The color scales are identical to that of Figure 5.2.*

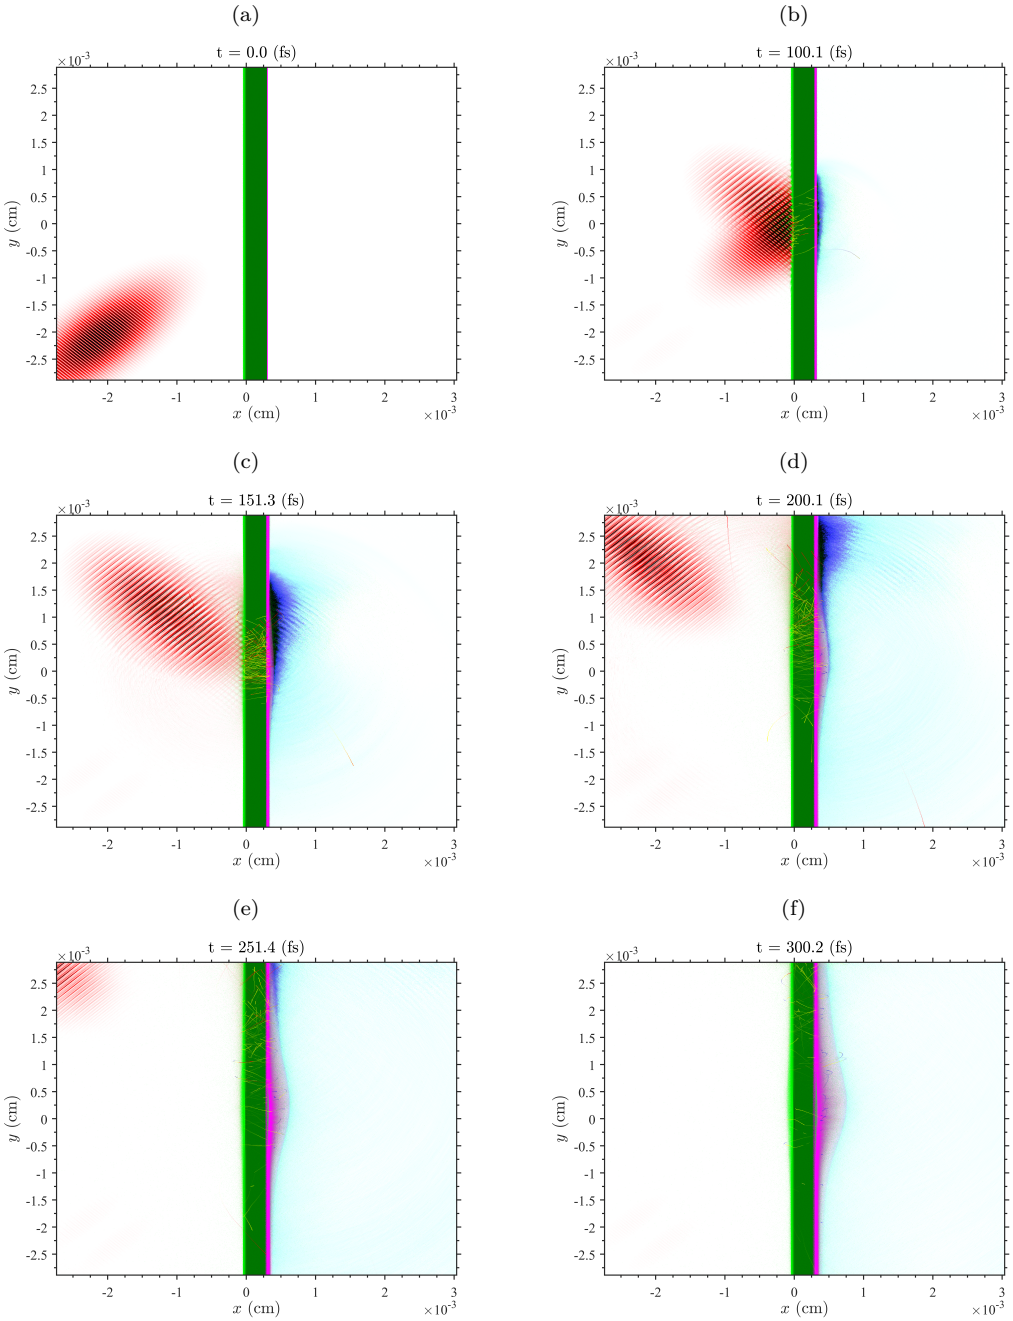

Figure 5.4: *The two figures show the proton energy spectra, corresponding to the simulation presented in Figure 5.3.*

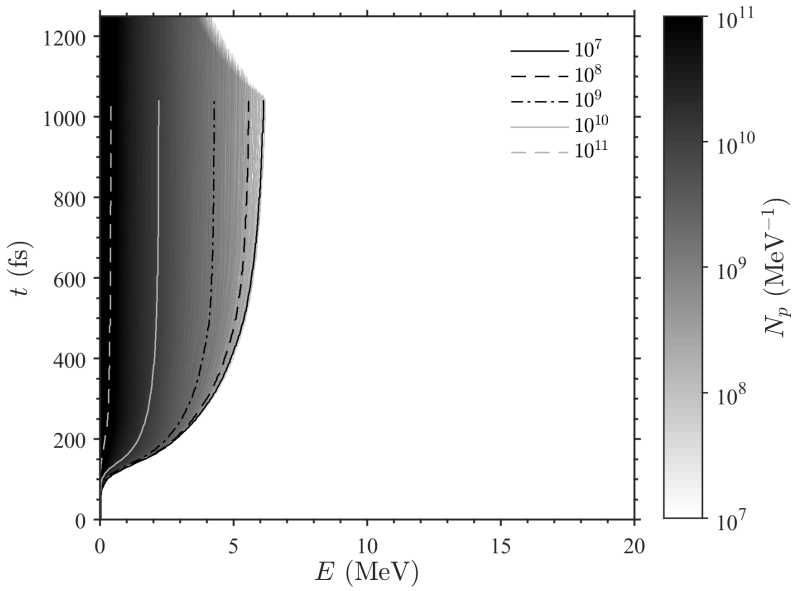

(a) *Time evolution of the proton energy spectrum. The lines correspond to the cutoff energy, E_c, and are calculated using different particle numbers, N_c, as described in Figure 4.1 and Equation* (4.4).

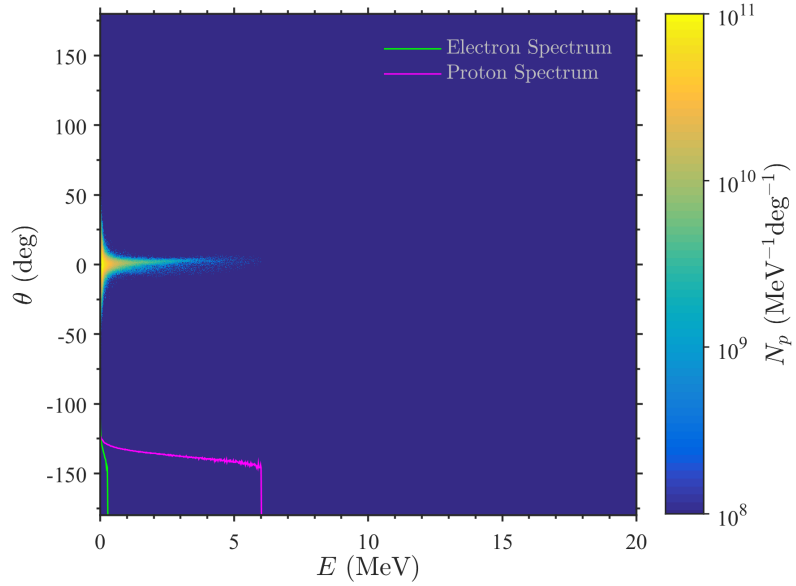

(b) *Angular distribution, with respect to traveling direction and relative to the target normal, of the proton energy spectrum, at 751.8 fs.*

5.1 Reducing the amount of reflected radiation

The first alternative target design that we will consider is aimed at reducing the amount of energy lost through the partial reflection of the pulse as it impinges on the front surface of the target. That is, we are here focusing on the very first stage presented in Figure 5.1. This has already been investigated in a number of different papers,[28,29,30,31] but will here be studied using nano-sized cones, which can be deposited on the surface of the target using so-called hole-mask colloidal lithography. These cones can easily be created with a base diameter of 50-500 nm, typically 2-3 diameters apart and with an angle of 20 degrees to the cone axis.[32]

Consequently, the here presented target design is an otherwise plain target but with cones embedded on its surface. The electron and ion densities in the cones are taken to be identical to that of the foil and the geometry of the cones are defined by a base diameter of 500 nm and an angle of 20 degrees. The cones are set to be uniformly separated in the transverse direction with a peak to peak distance of 1250 nm, corresponding to 2.5 base diameters. However, as the simulation is still performed in only two spatial dimensions it would be more correct to say that the target has gratings on the front surface, rather than cones, but given the size of the structures no major difference is expected and this colloquial denomination is therefore kept throughout the thesis.

Due to the presence of the cones, the treatment of the preplasma will have to be considered in a slightly different manner compared to previously. We still use the simple quadratically increasing density profile and the preplasma is still taken to start at a fixed distance from the foil surface, excluding the nano cones. This fixed distance is now taken to be the previous preplasma thickness plus the height of the cones. The preplasma density then quadratically increase to the foil density as it reaches the foil, including the cones. This way we attempt to, in a very simple and ad hoc manner, emulate a preplasma that has expanded from both the cones and the foil, partially filling the troughs.

Presented in Figure 5.5 are the electric field, the electron density and the proton density for different times into the simulation. As was also the case for the standard plain target, Figure 5.3, the charge-separation field has begun to form by the time the point of maximum intensity of the pulse has reached the target, Figure 5.5b. However, a closer comparison with the plain target shows that the field is both stronger and covers a larger spatial region than before. Furthermore, the charge-separation field now not only travels in the positive, but also the negative transverse direction and appears to remain for an extended amount of time as compared to the plain target, Figures 5.3c–5.3f. The reason for this appears to be twofold. The cones both increase the amount of energy transferred to the hot electrons and, as the target is no longer homogeneous in the y-direction, the directionality of the generated hot electrons has been altered such that the transverse electron transport is reduced.

The results of using the cones can be clearly seen by studying the time evolution of the proton energy spectrum, which for the setup at hand can be found in Figure 5.7a. The spectrum can be seen to be close to, but not entirely, saturated before the most highly energetic protons start leaving the simulation box. Before this point, the maximum energy achieved is three times higher than the equivalent of the plain target, Figure 5.3a. Upon comparing the angular distributions of the two setups, Figures 5.3a and 5.5a, we see that this however has come at the cost of now having a slightly less collimated proton beam.

Figure 5.5: *Simulation of a TNSA setup, with an pulse of 0.75 J incoming at an angle of 45 degrees to the target normal. The target consists of a 3 µm thick foil of plasma, with nano-sized cones on the front side, and is initially described by a density profile as shown in Figure 2.3. The preplasma extends from the foil such that the plasma density becomes zero at a fixed value of x. The color scales are identical to that of Figure 5.2.*

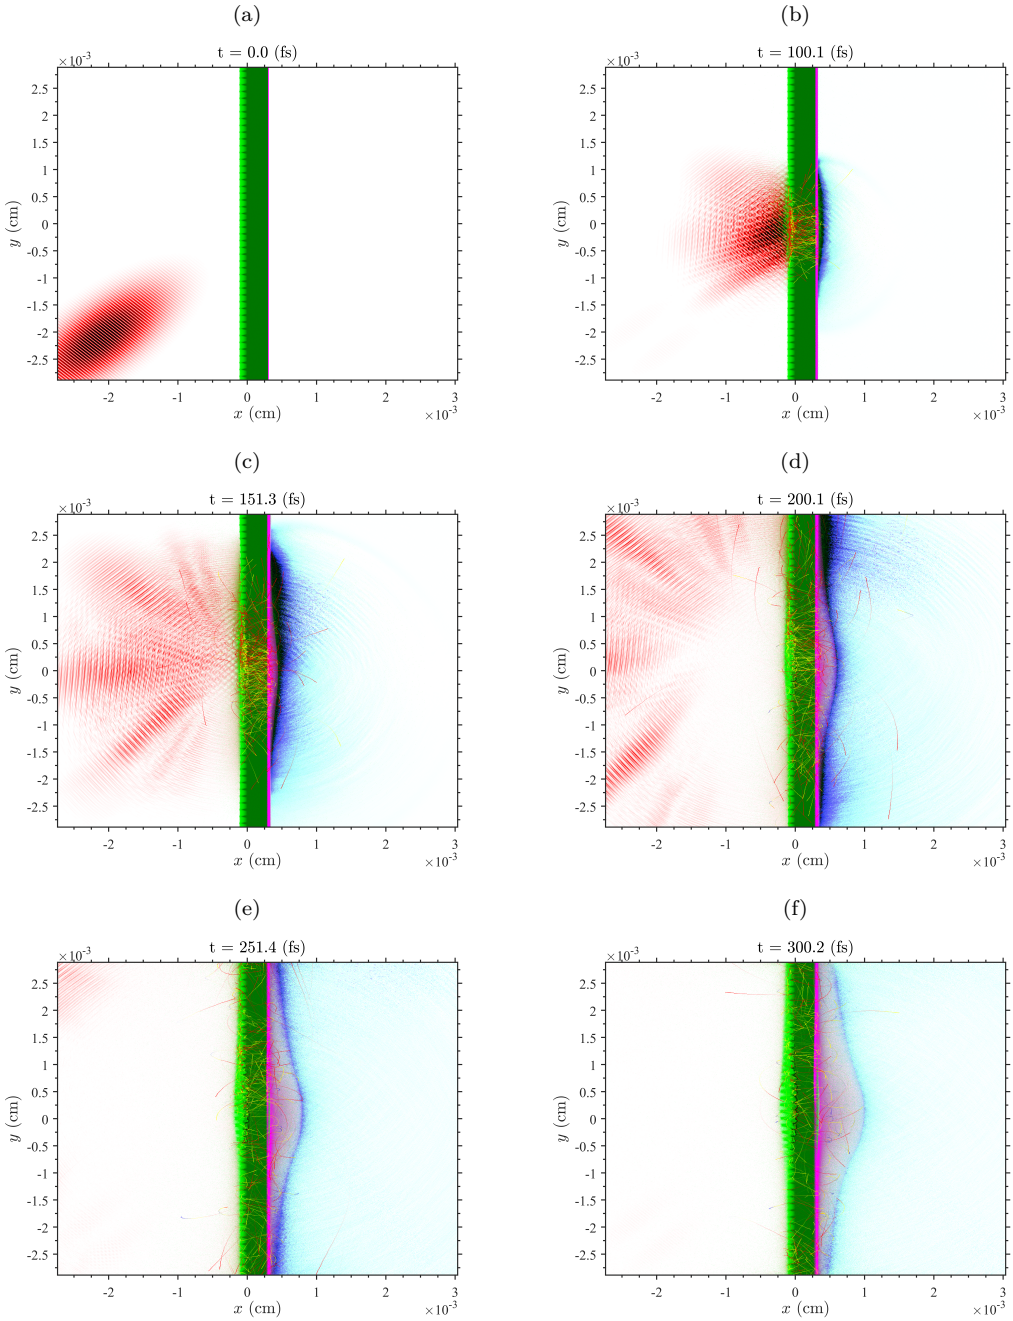

Figure 5.6: *The two figures show the proton energy spectra, corresponding to the simulation presented in Figure 5.5.*

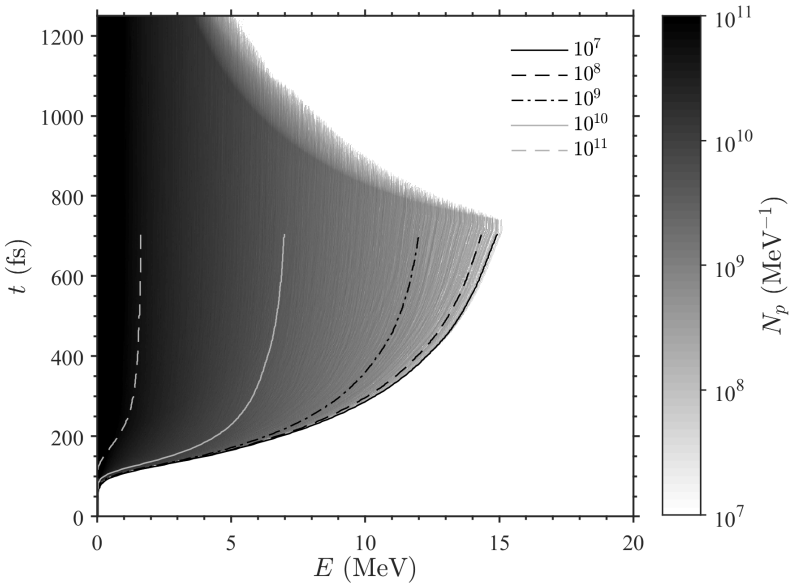

(a) *Time evolution of the proton energy spectrum. The lines correspond to the cutoff energy, E_c, and are computed using different particle numbers, N_c, as described in Figure 4.1 and Equation* (4.4).

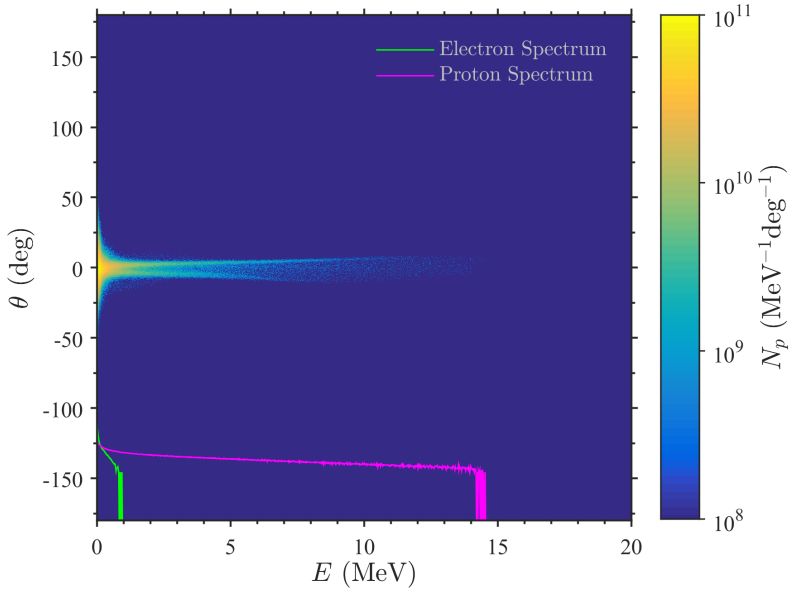

(b) *Angular distribution, with respect to traveling direction and relative to the target normal, of the proton energy spectrum, at 751.8 fs.*

5.2 Restricting hot electrons from transverse spreading

As was mentioned in the preceding section the transport of hot electrons in the transverse direction, the second stage presented in Figure 5.1, plays a major role in attempts at sustaining the charge-separation field for longer times. That this is indeed the case will be shown more explicitly in this section where we present two different setups that isolate the target in the transverse direction.

The first target towards this goal is a plain target, as in the standard TNSA setup. However, the target only extends 7.5 μm in the positive transverse direction, thus being equivalent of firing the laser pulse at the edge of the foil. The electric field, electron density and proton density for a few different times into the simulation are for the present setup presented in Figure 5.7. The effect of this, as compared to the plain, standard TNSA setup presented in Figure 5.3, is that hot electrons will travel in the positive transverse direction until they reach the edge of the foil, at which time they will be reflected at the plasma-vacuum boundary and instead escape the simulation region in the negative transverse direction. This also makes the transversely traveling charge-separation field get reflected, as can be seen in Figures 5.7b–5.7e.

Turning our attention to the proton energy spectrum generated using this setup, Figure 5.8, we see that this has increased the maximum proton energy by about 50 %, Figure 5.9a, as compared to the standard setup, Figure 5.5a. That we see this increase in energy is not a surprise. In fact, it is exactly what we envisioned would be the effect, and the very reason the setup was proposed. However, one may also suspect that there will be a change in the angular distribution, since the protons in the sheath mainly travel in direction normally to the target surface. That we will have protons traveling at 90 degrees to the target normal, from the edge where we also have envisioned a proton layer, is therefore a given. From Figure 5.9b we can see exactly this, also explaining why we now, in Figure 5.9a, see more than one occurrence of large amounts of particles leaving the simulation region. The first instance is due to the particles traveling in the transverse direction. Of greater importance is however the fact that the width of the main distribution, centered close to the target normal, appear to be unchanged, thus giving us the result of a more energetic proton beam without any effect on the collimation.

The second target presented in this section is also a plain target as the first one, but this time the target only extends 7.5 μm in both the positive and the negative transverse directions. This corresponds to having a thin, in both the normal and transverse direction, strip as the target still has an infinite extension in the third spatial dimension. As before we begin by studying the electric field, electron density and proton density for a few different times into the simulation, for this target design to be found in Figure 5.9.

From experience drawn from the first setup of this section, with the edge of the target, it is reasonable to suspect that isolating the target in also the negative transverse direction will lead to an even greater increase in energy, as the hot electrons should now be unable to be transported away from the region. This is also what we see from Figure 5.9, where the charge-separation field can be seen to be reflected twice over the time range covered, ultimately keeping a strong field present at the boundary between the expanding sheath and the vacuum for a longer time. By looking at the time evolution of the proton energy spectrum, Figure 5.11a, we find that the maximum proton energy has now increased by about 150 % as compared to the standard TNSA setup, Figure 5.5a. However, the increase is most likely even higher as the spectrum has not yet reached saturation by the time the protons start leaving the simulation box. As with the previous setup, the very first occurrence of particles leaving the region is in the transverse direction.

Furthermore, the angular distribution can be found in Figure 5.11b, showing particles traveling both in the positive and negative y-directions, as expected. The main distribution

this time hints at being slightly less collimated, thus giving us a slight trade-off for the large energy increase.

Figure 5.7: *Simulation of a TNSA setup, with an pulse of 0.75 J incoming at an angle of 45 degrees to the target normal. The target consists of a 3 µm thick plasma foil, which is initially described by a density profile as shown in Figure 2.3. The foil only extends 7.5 µ in the positive y-direction from the focus spot of the laser. The color scales are identical to that of Figure 5.2.*

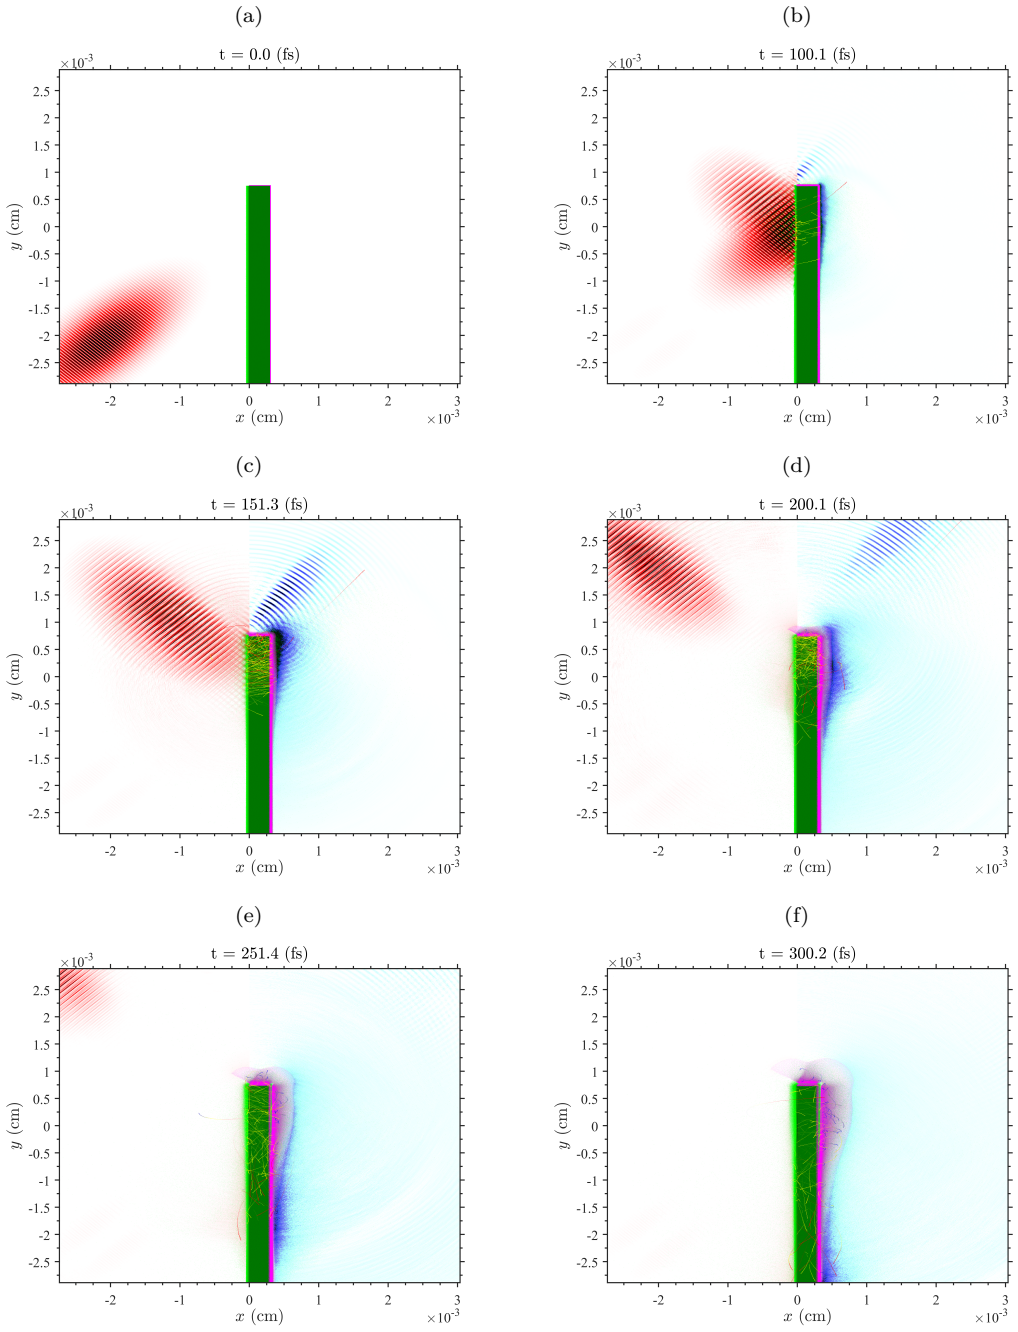

Figure 5.8: *The two figures show the proton energy spectra, corresponding to the simulation presented in Figure 5.7.*

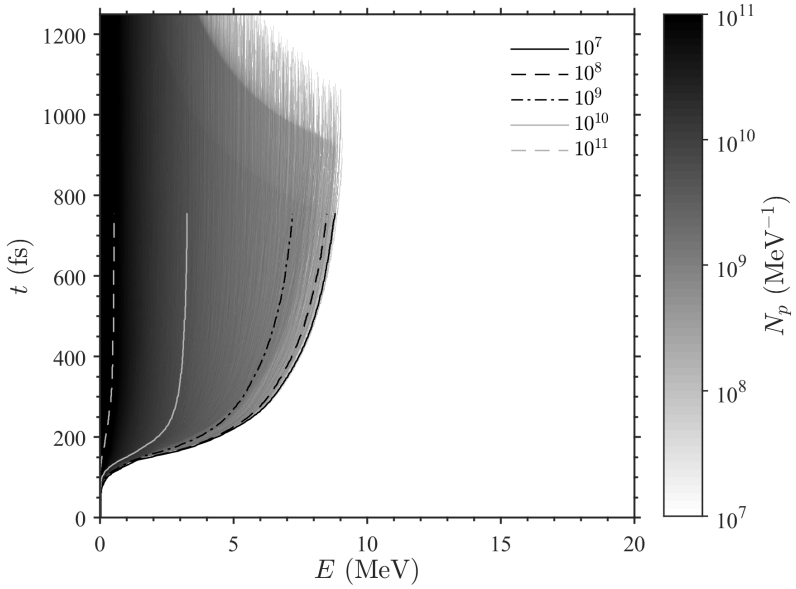

(a) *Time evolution of the proton energy spectrum. The lines correspond to the cutoff energy, E_c, and are calculated using different particle numbers, N_c, as described in Figure 4.1 and Equation* (4.4).

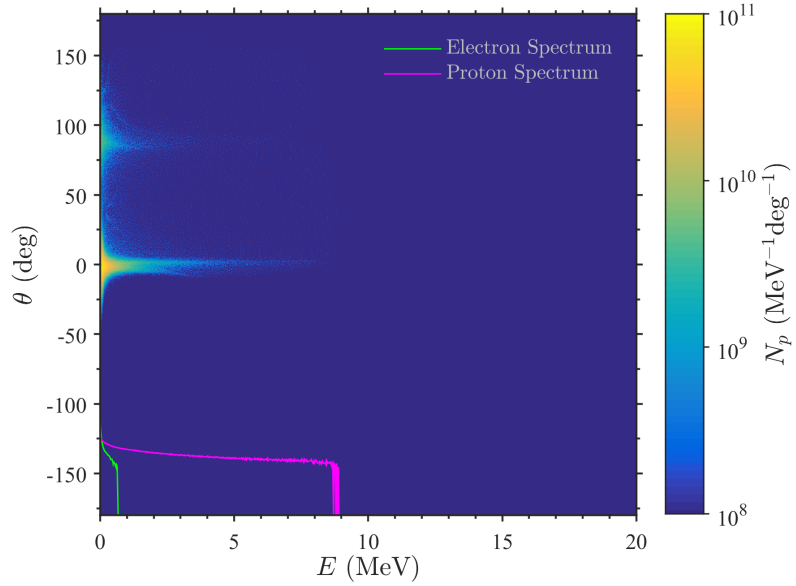

(b) *Angular distribution, with respect to traveling direction and relative to the target normal, of the proton energy spectrum, at 751.8 fs.*

Figure 5.9: *Simulation of a TNSA setup, with an pulse of 0.75 J incoming at an angle of 45 degrees to the target normal. The target consists of a block of plasma, of dimensions* $3\,\mu\text{m} \times 15\,\mu\text{m}$*, and is initially described by a density profile as shown in Figure 2.3. The color scales are identical to that of Figure 5.2.*

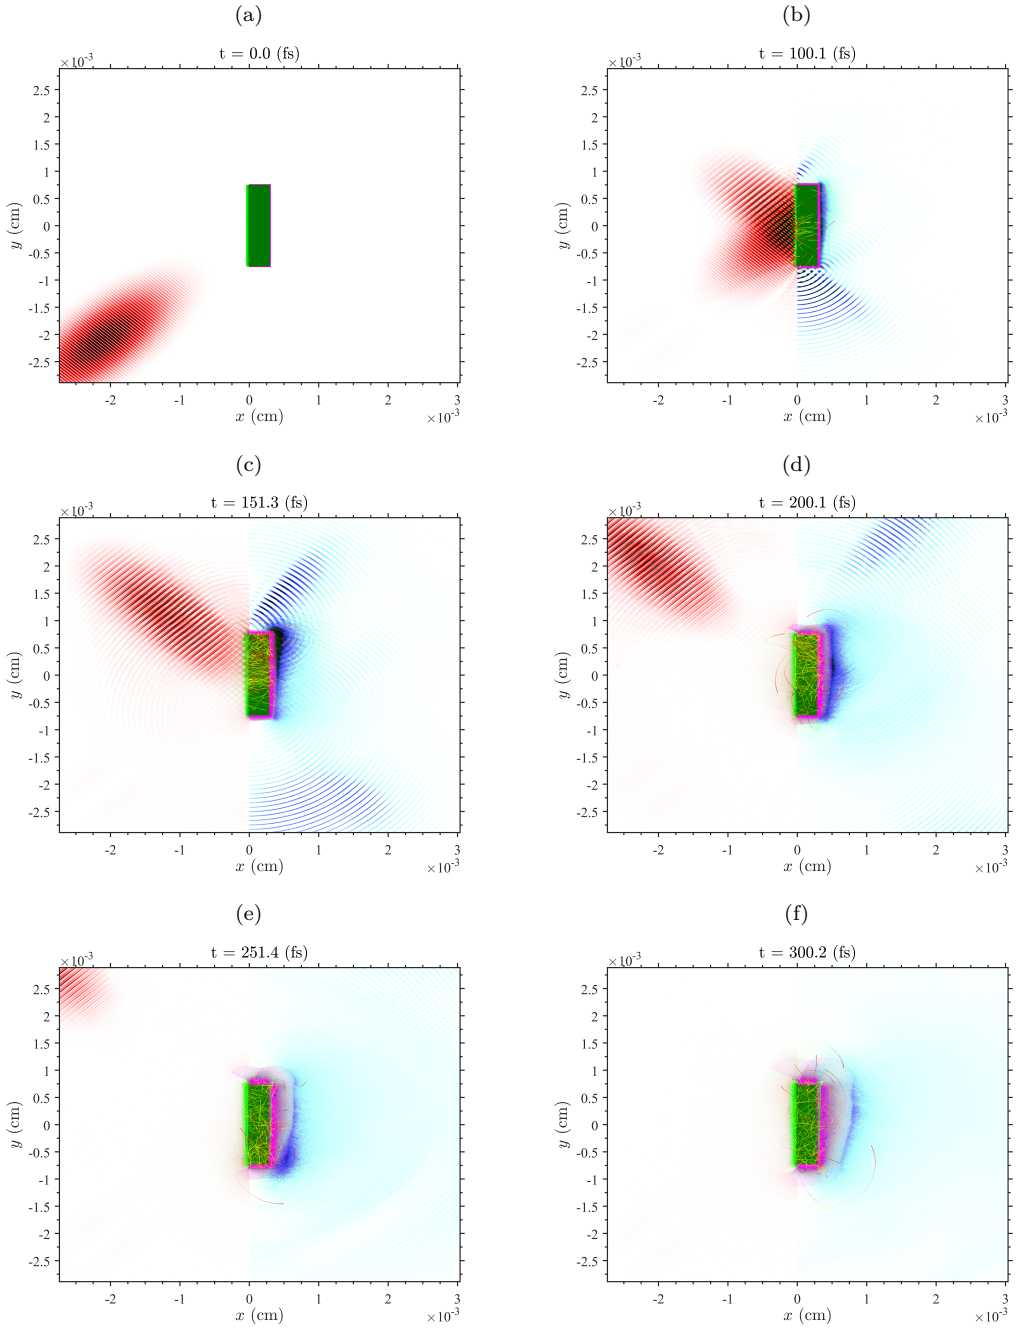

Figure 5.10: *The two figures show the proton energy spectra, corresponding to the simulation presented in Figure 5.9.*

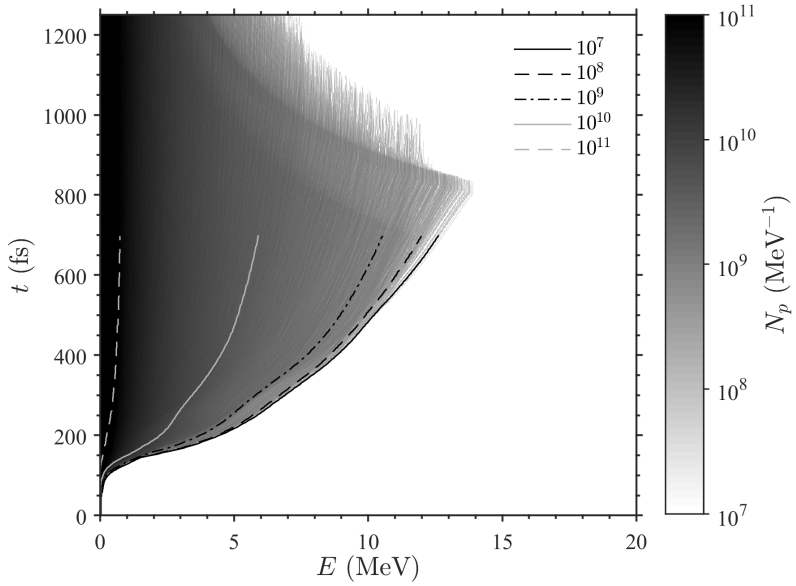

(a) *Time evolution of the proton energy spectrum. The lines correspond to the cutoff energy, E_c, and are calculated using different particle numbers, N_c, as described in Figure 4.1 and Equation* (4.4).

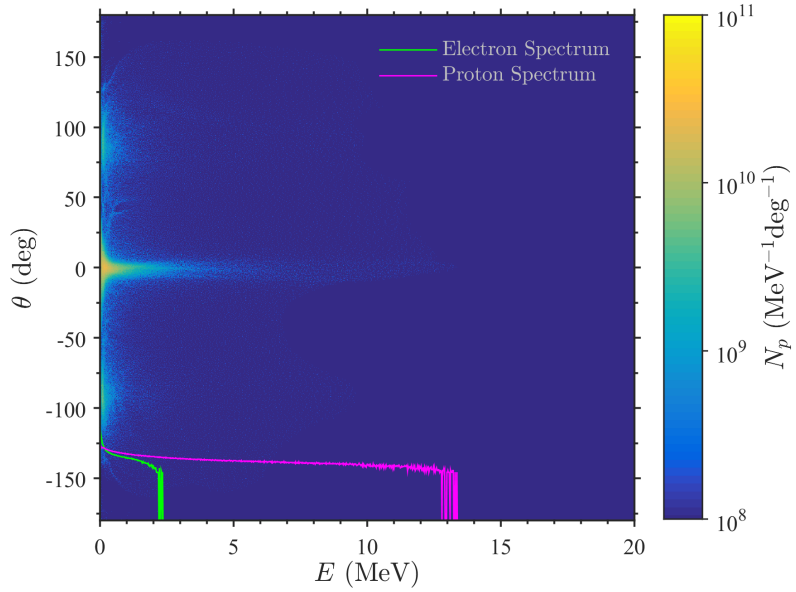

(b) *Angular distribution, with respect to traveling direction and relative to the target normal, of the proton energy spectrum, at 751.8 fs.*

5.3 Increasing the field at sharp tips

In the previous alternative setups we have seen how the charge-separation field can be increased and spatially localized to the region of interest for a longer time by means of increased absorption and restriction of the transverse movement of the highly energetic electrons. In this section we will present a final setup, drawing upon a rather different idea. It is well known that electric fields are stronger close to sharper objects.[18] A nice example of this was shown by Gonoskov et al.,[33] where the interaction of a laser pulse with a micron-sized Christmas tree was investigated. It is with this in mind that the final setup is utilizing nano-sized cones on the rear side of the target identical to what was used on the front side of the target in section 5.1. To reiterate, the electron and ion densities in the cones are identical to that of the foil, the geometry of the cones are defined by a base diameter of 500 nm and an angle of 20 degrees and the cones are set to be uniformly separated in the transverse direction with a peak to peak distance of 2.5 base diameters.

The presence of the cones on the rear surface requires us to alter the geometry of the proton sheath. The way we see it, this can be done in two ways. The easiest way is to simply let the sheath extend a fixed distance in the x-direction, from the surface (cones included). This would be equivalent of depositing an identical amount of protons per unit area on the rear surface. However, as the surface is no longer flat the sheath will be thinner at the sides of the cones, in the direction normal to the surface (cones included). The other alternative is then to make the proton sheath equally thick, with regards to the actual surface normal (cones included). This is presumably the more likely scenario of water vapour contamination.

We once again start by studying the electron and proton densities, along with the electric field, for a few different times into the simulation. This is presented in Figure 5.11, from which we can indeed see that the generated charge-separation field is stronger close to the tips of the cones. However, a comparison with the plain target, Figure 5.3, shows that the charge-separation field now covers a slightly smaller spatial region and appear to remain in the region on the same time scale as before. We can also see that it is mainly from the very tips of the cones that the expanding sheath of protons is forming and that these protons travel not only in the target normal direction as before, but in a wide range of directions due to the normal direction of the surface at the cones.

Upon turning towards the time evolution of the proton energy spectrum, Figure 5.13a, we find that the maximum proton energy has actually been slightly reduced compared to the standard setup, Figure 5.5a, contrary to the desired effect. The angular distribution, Figure 5.13b, however, shows what is expected. Namely that the collimation of the proton beam has all but vanished. This is simply an effect of the charge-separation field being directed normally to the rear target surface. As the rear surface now has inclined cones, the protons on the sides of the cones will be accelerated in the direction normal to this side rather than the x-direction as for a plain target.

The fact the cones failed to increase the proton energy is most likely due to the size of the structures. It can be seen that the structures on the rear side are quickly destroyed as the charge-separation field is formed. Thus, no noteworthy increase of the charge-separation field can be seen simply because the sharp edges do not survive long enough. Furthermore, the hot electrons appear to behave largely unaffected by the structures and are transported away just as easily as with the plain target. This could presumably be changed by increasing the size of the structures.

Figure 5.11: *Simulation of a TNSA setup, with an pulse of 0.75 J incoming at an angle of 45 degrees to the target normal. The target consists of a 3 µm thick foil of plasma, with nano-sized cones on the rear side, and is initially described by a density profile as shown in Figure 2.3. The plasma sheath on the rear side is extending a fixed distance in the positive x-direction. The color scales are identical to that of Figure 5.2.*

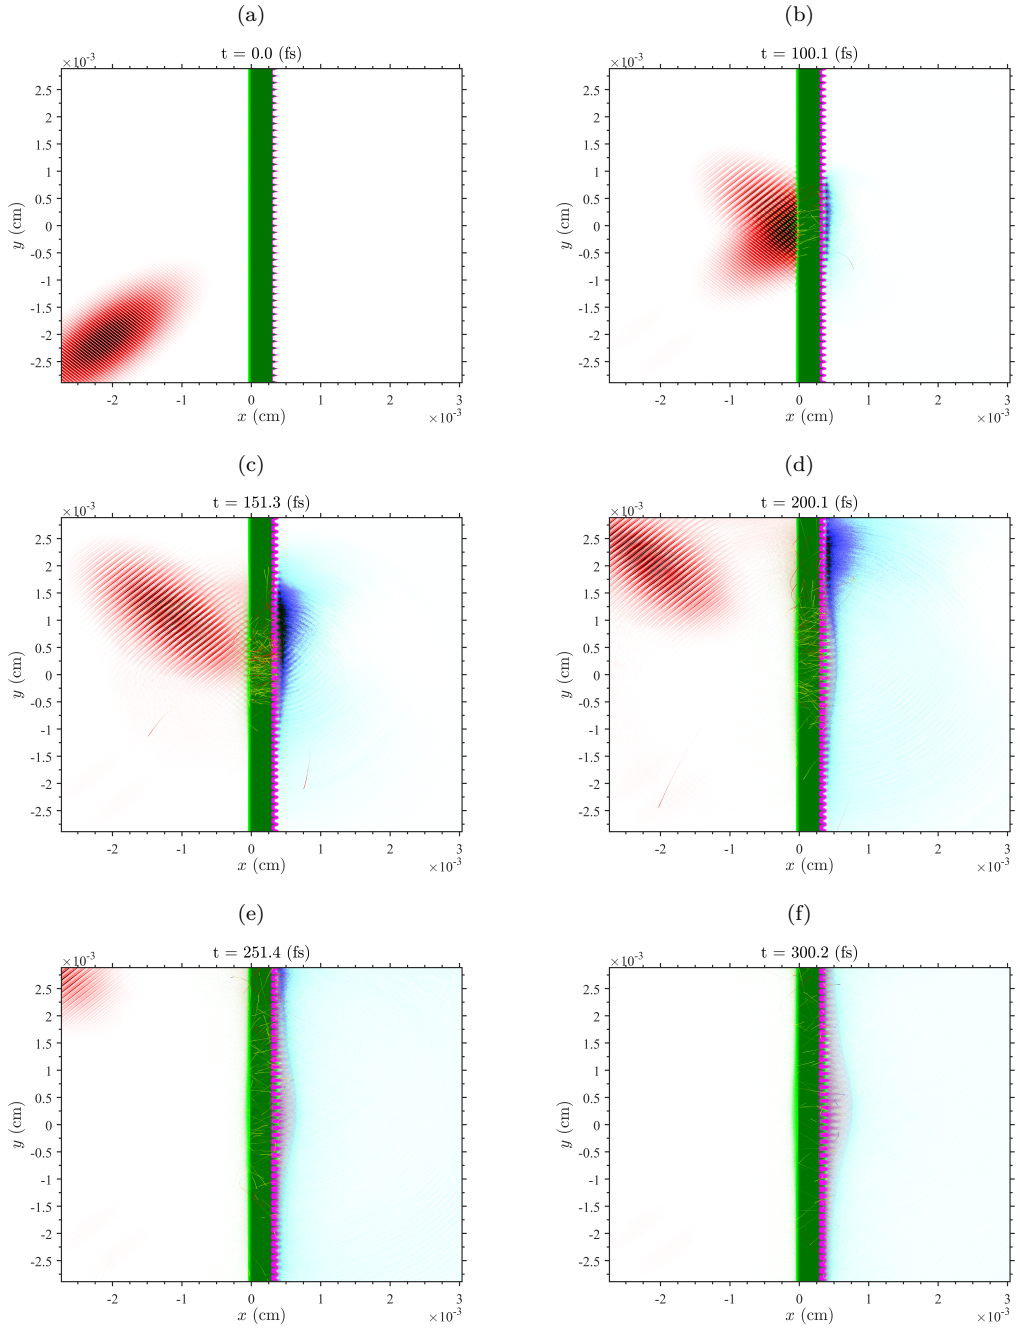

Figure 5.12: *The two figures show the proton energy spectra, corresponding to the simulation presented in Figure 5.11.*

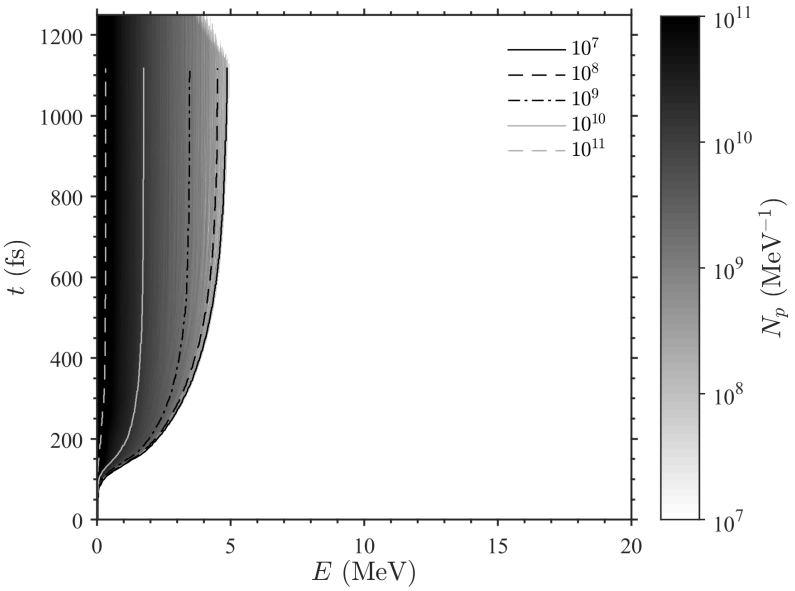

(a) *Time evolution of the proton energy spectrum. The lines correspond to the cutoff energy, E_c, and are computed using different particle numbers, N_c, as described in Figure 4.1 and Equation* (4.4).

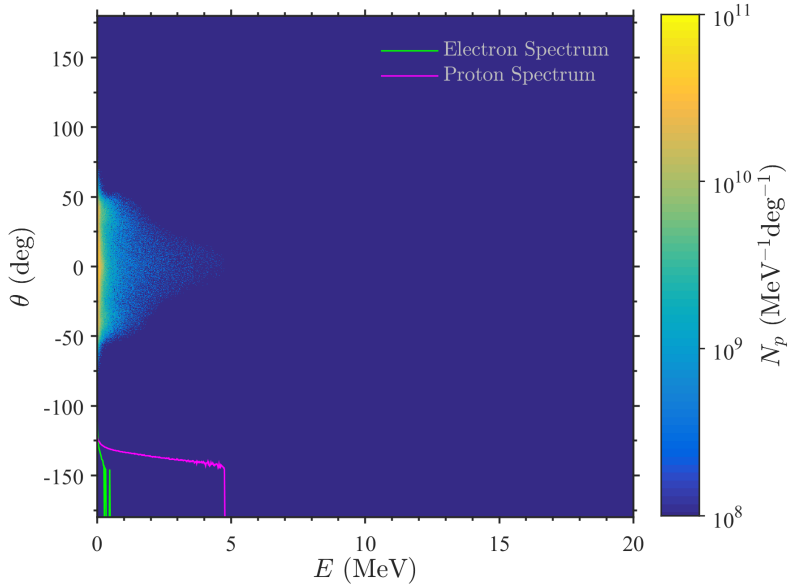

(b) *Angular distribution, with respect to traveling direction and relative to the target normal, of the proton energy spectrum, at 751.8 fs.*

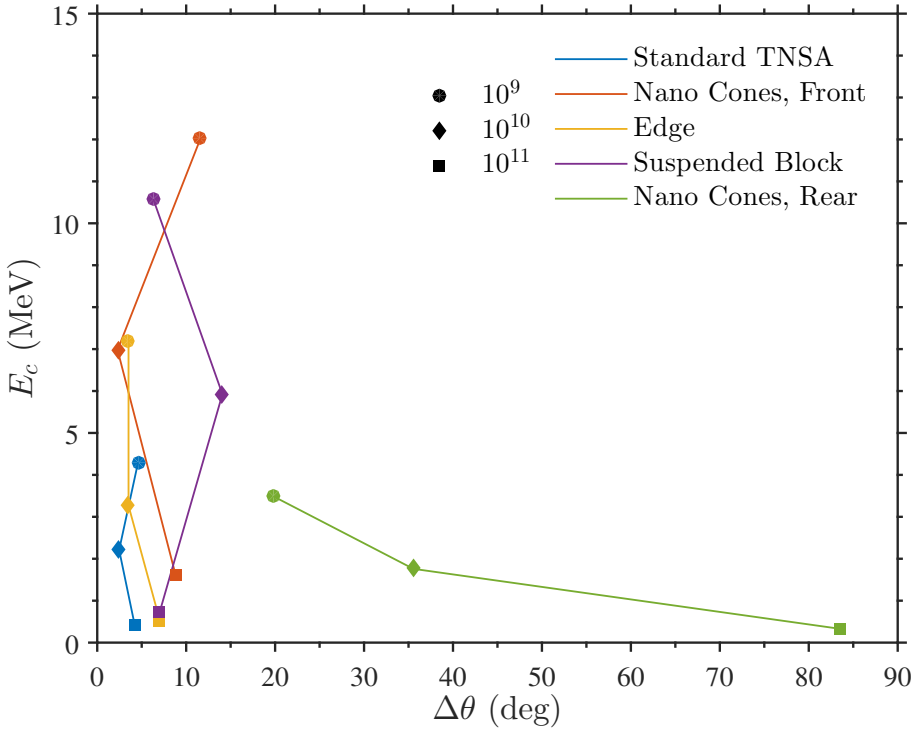

Figure 5.13: *Summary of the energy and collimation of the resulting proton beams from the simulation setups presented in Figures 5.3, 5.5, 5.7, 5.9 and 5.11. For every setup the energy, E_c, is computed for three different particle numbers, N_c, as described in Figure 4.1 and Equation* (4.4)*. Using only the part of the spectrum above this energy, the angular width, $\Delta\theta$, is taken as the value of full width at half maximum of the central peak in the corresponding angular spectrum.*

5.4 Summary and suggestions

In this chapter four alternative target designs have been described and analyzed according to proton energy spectra and angular distributions of direction of travel. These have mainly been compared to the standard TNSA setup, with a flat surface, presented in the beginning of the chapter and the basis for the target designs was mostly draw from ideas presented in Figure 5.1.

An attempt has been made at summarizing the performance of the different setups with respect to the parameters already mentioned, using the definition of the energy cutoff from Equation 4.4. This is presented in Figure 5.13 and for three different particle numbers, N_c. The angular width, $\Delta\theta$, is the value of full width at half maximum of the distribution centered at 0 degrees in the angular spectrum. Only the part of the angular spectrum above the corresponding energy cutoff is used in calculating the angular width.

What the figure shows is that there is only a mild change in the collimation of the proton beams between the different setups, but that the energy cutoff can be increased substantially by either increasing the absorbed energy or decreasing the transverse transport of electrons. It also clearly shows that the addition of nano-sized cones on the rear side of the target greatly affects the collimation of the proton beam, for the worse, without any increase in proton energy, at least for the size and geometry considered in this work.

Upon performing the simulation of the target limited in both the positive and negative y-direction it was noticed that the hot electrons circling around the target appear to do so quite isotropically. It has also been noticed that the currents generated within the target gives rise to strong magnetic fields in the near vicinity of the target surface. Through this, an idea came to mind that a very strong magnetic field could be used to restrict the movement of the hot electrons in the transverse directions, thereby making the charge separation field stronger and presumably more directed.

A suggestion is therefore to use one or multiple laser pulses to first set up a strong current in the target, such as to magnetize it. This can be done by having the target be a micron-sized puck which you hit from the side, but slightly off center. This will set up a strong circular current in the puck, generating a magnetic field in the target normal direction. When the main pulse hits the face of the puck the hot electrons will presumably behave much like in the simulation most recently mentioned, only this time their transverse motion should be more restricted and thereby improve the quality of the proton beam as discussed.

Chapter 6

Conclusion

We have in this thesis plunged into the field of ultra intense laser-matter interactions. This has mainly been done through the use of particle-in-cell simulations via the numerical code Picador, focused on light ion acceleration in the TNSA regime. The light ions have, for all intents and purposes of this thesis, been represented by protons.

The first question that we set out to answer was if TNSA simulations in Picador are affected by any artificial charging of the simulation boundary. It was shown in chapter 4 that this effect is insignificant for typical TNSA simulations, given that the simulation box is sufficiently large. Furthermore, it was noted that this box size is roughly the same as the size required for the proton beam to reach saturation in energy. For 2D simulations considered in this thesis it means that typical TNSA simulation can be performed with a simulation box on the order of $50\,\mu\text{m} \times 50\,\mu\text{m}$.

In chapter 4 we also identified and resolved a stability issue present in 2D (as well as in 1D) simulations when using Picador. This was caused by incorrect handling of current deposition and boundary conditions when the number of grid cells in at least one dimension became too small. Furthermore, another implementation error was concurrently found and fixed, this time relating to the CIC and TSC current deposition schemes. This issue came from the fact that particle positions and momenta are defined at different moments in time (differing by half a time step), something that was not properly handled. Both of these implementation errors were eventually spotted much through the use of monitoring the total energy in the simulation volume. The diagnostics tool to do this was implemented as a module in Picador, as part of the work on this thesis, and further improvements of the module have been suggested.

As a result of the work presented in chapter 4 it became possible to perform efficient and reliable 2D TNSA simulations using Picador. We have taken full advantage of this in chapter 5, where a number of alternative target designs have been investigated. The designs were chosen such as to address key ideas of how to improve the TNSA process. The effects of adding nano-sized cones to both the front and rear target surfaces have been studied in an attempt at increasing the amount of laser energy absorbed by the target and increasing the charge separation field, respectively. Furthermore, two target designs with a finite and semi-infinite extension in the (simulated) transverse dimension, respectively, have been investigated. The key idea with these designs is to reduce the unwanted transport of hot electrons in the transverse direction.

Upon comparing the discussed alternative designs to a more typical plain target it was found that all but one of the designs resulted in an improvement of the proton beam. The usage of nano-sized cones on the front surface greatly increased the proton energy, both by increasing the amount of laser radiation absorbed by the target but also by making the generating the hot electrons more isotropically, thus decreasing their transverse spreading. The two designs of a target with limited extension also increased the proton energy, without

much of a trade-off in collimation, in the way that was expected. Thus, it was the use of nano-cones on the rear surface that did not result in any improvement of the proton beam. The reason most likely being that the structures are too small to have any significant effect on the hot electrons.

Lastly, a new target design is proposed at the end of chapter 5 involving the magnetization of a micron-sized puck, using one or even multiple laser pulses. The generated magnetic field is then though to restrict the ability of the hot electrons, generated by a main pulse, to move in the transverse directions. This design is however not studied further in this thesis, but is left for future investigation.

References

[1] T. H. Maiman. Stimulated optical radiation in ruby. *Nature* **187** (Aug. 1960), 493–494. URL: http://dx.doi.org/10.1038/187493a0.

[2] G. A. Mourou, C. P. J. Barry, and M. D. Perry. Ultrahigh-Intensity Lasers: Physics of the Extreme on a Tabletop. *Phys. Today* (1998).

[3] V. L. Ginzburg. *Propagation of Electromagnetic Waves in Plasma*. Gordon & Breach Science Publishers Ltd, 1961.

[4] D. W. Forslund, J. M. Kindel, and K. Lee. Theory of Hot-Electron Spectra at High Laser Intensity. *Phys. Rev. Lett.* **39** (5 Aug. 1977), 284–288. DOI: 10.1103/PhysRevLett.39.284.

[5] K. Estabrook and W. L. Kruer. Properties of Resonantly Heated Electron Distributions. *Phys. Rev. Lett.* **40** (1 Jan. 1978), 42–45. URL: http://link.aps.org/doi/10.1103/PhysRevLett.40.42.

[6] D. W. Forslund and J. U. Brackbill. Magnetic-Field-Induced Surface Transport on Laser-Irradiated Foils. *Phys. Rev. Lett.* **48** (23 June 1982), 1614–1617. DOI: 10.1103/PhysRevLett.48.1614.

[7] V. I. Veksler. "Coherent principle of acceleration of charged particles". *Proc. CERN Symp. of High Energy Accelerators and Pion Physics*. Geneva, Switzerland, 1956.

[8] D. Strickland and G. Mourou. Compression of amplified chirped optical pulses. *Optics Communications* **56**.3 (1985), 219–221. DOI: http://dx.doi.org/10.1016/0030-4018(85)90120-8.

[9] P. F. Moulton. Ti-doped sapphire: tunable solid state laser. *Optics News* **8** (1982), 9.

[10] P. F. Moulton. Spectroscopic and laser characteristics of Ti:Al2O3. *J. Opt. Soc. Am. B* **3**.1 (Jan. 1986), 125–133. DOI: 10.1364/JOSAB.3.000125.

[11] C. Danson et al. Vulcan Petawatt—an ultra-high-intensity interaction facility. *Nuclear Fusion* **44**.12 (2004), S239. URL: http://stacks.iop.org/0029-5515/44/i=12/a=S15.

[12] Berkeley Lab News Center. *BELLA: The Next Stage in Laser Wakefield Acceleration*. Aug. 2015. URL: http://www2.lbl.gov/publicinfo/newscenter/features/2008/apr/af-bella.html.

[13] The official ELI website. Aug. 2015. URL: http://www.eli-laser.eu/.

[14] Lund High-Power Laser Facility. Aug. 2015. URL: http://www.atomic.physics.lu.se/laser-facility/.

[15] A. Gonoskov. "Ultra-intense laser-plasma interaction for applied and fundamental physics". PhD thesis. Umeå University, 2014.

[16] C. K. Birdsall and A. B. Langdon. *Plasma Physics via Computational Simulation*. McGraw-Hill, 1985.

[17] D. Tskhakaya et al. The Particle-In-Cell Method. *Contributions to Plasma Physics* **47**.8-9 (2007), 563–594. DOI: 10.1002/ctpp.200710072.

[18] J. D. Jackson. *Classical Electrodynamics*. 3rd ed. Wiley, 1998.

[19] H. Okuda. Nonphysical noises and instabilities in plasma simulation due to a spatial grid. *Journal of Computational Physics* **10**.3 (1972), 475–486. DOI: http://dx.doi.org/10.1016/0021-9991(72)90048-4.

[20] S. Bastrakov et al. Particle-in-cell plasma simulation on heterogeneous cluster systems. *Journal of Computational Science* **3**.6 (2012). Next Generation Computational Scientists: Russian Federation, 474–479. DOI: http://dx.doi.org/10.1016/j.jocs.2012.08.012.

[21] Website of the PICADOR project. Aug. 2015. URL: http://hpc-education.unn.ru/en/research/overview/laser-plasma.

[22] I. Surmin et al. "Dynamic Load Balancing Based on Rectilinear Partitioning in Particle-in-Cell Plasma Simulation". *Parallel Computing Technologies*. Vol. 9251. Lecture Notes in Computer Science. Springer International Publishing, 2015, pp. 107–119. DOI: 10.1007/978-3-319-21909-7_12.

[23] T. Rylander, A. Bondeson, and P. Ingelström. *Computational Electromagnetics*. 2nd ed. Springer, 2012.

[24] M. Burza et al. Hollow microspheres as targets for staged laser-driven proton acceleration. *New Journal of Physics* **13**.1 (2011), 013030. URL: http://stacks.iop.org/1367-2630/13/i=1/a=013030.

[25] A. Macchi, M. Borghesi, and M. Passoni. Ion acceleration by superintense laser-plasma interaction. *Rev. Mod. Phys.* **85** (2 May 2013), 751–793. DOI: 10.1103/RevModPhys.85.751.

[26] The group for Ultra-High Intensity Laser Physics. May 2015. URL: http://www-atom.fysik.lth.se/Ultrahigh/index.htm.

[27] A. Gonoskov. "Relativistic effects in the problems of charged particle acceleration and generation of gamma and X-ray emission in the interaction of femtosecond laser pulses with plasma structures". Russian. PhD thesis. 2011.

[28] Y. Nodera et al. Improvement of energy-conversion efficiency from laser to proton beam in a laser-foil interaction. *Phys. Rev. E* **78** (4 Oct. 2008), 046401. DOI: 10.1103/PhysRevE.78.046401.

[29] L. Cao et al. Enhanced absorption of intense short-pulse laser light by subwavelength nanolayered target. *Physics of Plasmas* **17**.4, 043103 (2010). DOI: http://dx.doi.org/10.1063/1.3360298.

[30] Z. Zhao et al. Acceleration and guiding of fast electrons by a nanobrush target. *Physics of Plasmas* **17**.12, 123108 (2010). DOI: http://dx.doi.org/10.1063/1.3507292.

[31] D. Margarone et al. Laser-Driven Proton Acceleration Enhancement by Nanostructured Foils. *Phys. Rev. Lett.* **109** (23 Dec. 2012), 234801. DOI: 10.1103/PhysRevLett.109.234801.

[32] H. Fredriksson et al. Hole–Mask Colloidal Lithography. *Advanced Materials* **19**.23 (2007), 4297–4302. DOI: 10.1002/adma.200700680.

[33] A. Gonoskov et al. Lighting up the Christmas tree: high-intensity laser interactions with a nano-structured target (2014). URL: http://arxiv.org/abs/1412.5998.