


POLITECNICO DI MILANO

**Scuola di Ingegneria Industriale e dell'Informazione
Corso di Laurea in Ingegneria Matematica**



**Intense Laser Interaction with
Nanostructured Plasmas:
a Kinetic Numerical Investigation**

Supervisor: Prof. Matteo Passoni

Co-supervisor: PhD. Luca Fedeli

**Dissertation of:
Arianna Formenti
Personal ID 823972**

Academic Year 2015-2016

Abstract

This thesis reports on a kinetic numerical investigation of the interaction between high intensity laser pulses and nanostructured plasmas in a regime characterized by the physical process of ion acceleration. A laser pulse is irradiated onto a double-layer target consisting in a low-density, nanostructured foam attached to a thin solid foil, giving rise to an enhanced acceleration process. The particle-in-cell code *piccante* is employed to solve the dynamics of the interaction, whose governing equations are given by the relativistic Vlasov-Maxwell system. Large scale two-dimensional and three-dimensional simulations, representative to some degree of the experiments, are performed to gain insights on the physical mechanisms at play. Special attention has been addressed towards the mathematical modelling of the foam morphology and structure, taking into account density and thickness non-uniformities, so to acquire elements concerning the role of the nanostructure in the interaction process. Moreover, plausible simulation parameters have been chosen for comparisons with available experimental data. The will of performing larger sets of simulations with setups closer to the experimental ones emerged, once ascertained feasibility and limits in the use of realistic values for the physical system parameters. This led to the necessity of improving the code with further optimizations by developing the OpenMP parallelization model into the code.

Sommario

Questa tesi riporta un'indagine numerica sull'interazione tra impulsi laser ad alta intensità e plasmi nanostrutturati in un regime caratterizzato dal processo fisico di accelerazione di ioni. Si considera l'irraggiamento da parte di un impulso laser su un bersaglio doppio strato composto da un film nanostrutturato a bassa densità, detto schiuma, depositato su di un sottile foglio solido. La dinamica cinetica e relativistica dell'interazione, governata dal sistema di equazioni di Vlasov-Maxwell, viene risolta numericamente mediante l'utilizzo del codice particle-in-cell *piccante*. Con l'obiettivo di analizzare i meccanismi fisici in gioco, il sistema in esame viene riprodotto attraverso onerose simulazioni in due e tre dimensioni, in parte rappresentative degli esperimenti. La modellazione matematica della morfologia e della struttura delle schiume è stata considerata con particolare attenzione, più specificamente tenendo conto delle disomogeneità nella densità e nello spessore, in modo da ottenere elementi sul ruolo della nanostruttura nella fisica dell'interazione. Le simulazioni sono state predisposte con parametri verosimili, così da poter effettuare confronti con i dati sperimentali disponibili in letteratura. Dopo aver accertato la possibilità e i limiti nell'utilizzo di valori realistici per i parametri fisici del sistema, è maturato l'interesse per un maggior numero di simulazioni con setup simili a quelli sperimentali. Ciò ha condotto alla necessità di introdurre ulteriori ottimizzazioni all'interno del codice, in particolare sviluppando il modello di programmazione parallela OpenMP.

Estratto

Dalla sua invenzione nel 1960 sino ad oggi, lo sviluppo della tecnologia laser ha reso possibili innumerevoli applicazioni scientifiche e tecnologiche. Questo lavoro di tesi si inserisce nell'ambito dello studio dell'interazione laser-materia in un regime caratterizzato da altissime intensità ($I > 10^{18}$ W/cm²). Poiché già a partire da intensità molto inferiori qualunque materiale viene ionizzato in pochi cicli di laser, l'interazione avviene in ogni caso con materia nello stato di plasma.

Diverse tecnologie laser all'avanguardia possono generare impulsi ultrabrevi (durata di decine di fs) ad alta potenza (~ 1 PW), focalizzati su aree molto ridotte (qualche μm^2 , vicino al limite di diffrazione), il che comporta intensità estremamente elevate ($\gtrsim 10^{21}$ W/cm²). Tali impulsi si dicono *superintensi*. L'attività di ricerca sull'interazione laser-plasma in questi regimi negli ultimi decenni ha portato ad individuare svariati schemi di potenziale interesse per varie applicazioni, tra cui l'accelerazione di particelle. Fra gli schemi di accelerazione di particelle, l'accelerazione di ioni mediante plasmi prodotti da laser risulta di particolare interesse. Schemi di questo tipo potrebbero consentire in futuro la realizzazione di sorgenti compatte di ioni con energie superiori a 100 MeV, con caratteristiche tali da renderli preferibili ad altre tecniche tradizionali di accelerazione. Tuttavia, attualmente diversi ostacoli devono essere superati per rendere le sorgenti di ioni "laser-driven" adatte per le applicazioni. Questo lavoro di tesi si inserisce in un progetto di ricerca finalizzato a migliorare le proprietà di sorgenti di ioni "laser-driven" con lo scopo di avvicinarle ai requisiti richiesti da alcune delle applicazioni summenzionate.

Uno degli schemi più studiati per l'accelerazione di ioni con plasmi prodotti da laser prevede di irraggiare bersagli solidi sottili con laser superintensi. A seguito dell'interazione, una frazione degli elettroni del bersaglio acquisisce un'elevata energia cinetica tale da permettergli di attraversare il bersaglio e generare sul suo retro una nuvola elettronica. Ciò genera una forte separazione di carica, ovvero un campo elettrico longitudinale molto intenso che a sua volta porta all'accelerazione degli ioni del bersaglio. I primi ad essere accelerati sono gli ioni H e C che vengono dai contaminanti che si trovano sul retro del bersaglio. I fasci di ioni accelerati che si ottengono con laser a fs sono caratterizzati da energie massime ~ 10 MeV, largo spettro esponenziale, con un numero particelle accelerate $\sim 10^{10}$.

Diverse strategie sono state proposte per superare alcuni limiti dell'accelerazione di ioni con laser. Una strategia emersa negli ultimi anni si basa sull'ottimizzazione del suddetto schema, in particolare aumentando l'efficienza dell'accoppiamento fra l'impulso laser incidente e gli elettroni del bersaglio, con lo scopo di aumentarne la temperatura (ovvero l'energia cinetica) e quindi l'entità del campo elettrico accelerante. Per farlo è possibile ingegnerizzare opportunamente i bersagli convenzionalmente utilizzati negli esperimenti di accelerazione di ioni, ovvero fogli solidi sottili, progettando nuove configurazioni di bersagli avanzati. Tra questi vi sono particolari bersagli doppio strato composti da un materiale nanostrutturato a bassa densità (qualche mg/cm^3 , che corrisponde alla densità critica per un laser superintenso generato per esempio da un sistema Titanio:Zaffiro), detto *schiuma*, depositato su di un foglio solido sottile. Le schiume sono concepite con l'idea di incrementare l'assorbimento di energia del laser da parte del bersaglio irraggiato sfruttando il forte accoppiamento che si instaura nell'interazione tra un plasma a densità critica (o quasi) e un tipico impulso laser superintenso. In ultima istanza, ciò porta ad un miglioramento delle prestazioni del processo di accelerazione. A causa dei bassi valori di densità richiesti insieme alla necessità di far aderire la schiuma al foglio solido, la produzione di tali bersagli è tutt'altro che banale. Una delle poche tecniche disponibili è la *pulsed laser deposition*, la quale permette di crescere lo strato di schiuma direttamente sul substrato solido. I meccanismi in gioco durante il processo di crescita portano inevitabilmente ad ottenere materiali porosi e nanostrutturati, ovvero costituiti da aggregati di nanoparticelle (gruppi di particelle di dimensione dell'ordine di 10 nm) che si assemblano in particolari strutture sulla mesoscala (scale $\gtrsim \mu\text{m}$).

Molti complessi meccanismi fisici caratterizzano l'interazione tra impulsi laser superintensi e materia. Ad ogni modo, la materia viene fortemente ionizzata, diventando plasma. Ciò significa che può essere considerata come composta da una collezione di particelle cariche che interagiscono tra di loro attraverso i campi elettromagnetici generati dalla loro stessa dinamica. Nell'ambito dell'accelerazione di ioni, le grandezze caratteristiche sono tali da poter considerare l'interazione elettromagnetica a lungo raggio dominante su quella a corto raggio, cosicché il comportamento del plasma è prevalentemente collettivo. Gli studi di carattere teorico relativi a questo tipo di sistemi si fondano principalmente sulla teoria cinetica dei plasmi. In particolare, il modello matematico più adatto è dato dall'equazione di Vlasov relativistica accoppiata alle equazioni di Maxwell. Studi analitici di questo sistema di equazioni sono possibili solo sotto opportune ipotesi semplificatrici e non si prestano all'analisi dettagliata di configurazioni reali come quelle di interesse nel contesto dell'accelerazione di ioni da laser. In questo caso occorre adottare degli opportuni schemi numerici, tra i quali il metodo *particle-in-cell* (PIC) risulta essere il più adatto per le simulazioni cinetiche dell'interazione laser-plasma in regime relativistico, oltre che il più largamente utilizzato.

L'obiettivo di questo lavoro consiste in un'indagine numerica dell'accelerazione di ioni

mediante l'interazione di impulsi laser superintensi ed ultra-corti con plasmi nanostrutturati nel cosiddetto regime *near-critical* (legato alla bassa densità della schiuma) in cui si manifesta una grande varietà di meccanismi complessi. Il proposito generale è quello di esplorare la possibilità di utilizzare lo strumento PIC per lo studio della fisica di tale interazione mediante onerose simulazioni in tre e due dimensioni, in parte rappresentative degli esperimenti. La modellazione della morfologia e della struttura del bersaglio è stata considerata con particolare attenzione con il fine di ottenere elementi sul ruolo della schiuma nell'interazione. Alcune proprietà della schiuma sono state fatte variare, mentre i parametri delle simulazioni sono stati scelti confrontabili con quelli di alcuni esperimenti descritti in letteratura. Dopo aver accertato la possibilità e i limiti nell'utilizzo di valori realistici per i parametri fisici del sistema, è maturato l'interesse per effettuare un maggior numero di simulazioni con setup quanto più simili a quelli sperimentali. Ciò ha condotto alla necessità di introdurre ulteriori ottimizzazioni all'interno del codice, in particolare sviluppando il modello di programmazione parallela OpenMP.

Nelle simulazioni tridimensionali le schiume nanostrutturate sono ottenute con il modello *diffusion-limited aggregation*, che genera schiume con morfologie e strutture piuttosto simili a quelle reali, così da tenere in conto le particolari disomogeneità nella densità che caratterizzano le schiume. Sono state poi considerate le disomogeneità nello spessore medio per indagare il loro effetto sull'accelerazione degli ioni. Le schiume nanostrutturate bidimensionali, invece, sono costituite da una collezione di nanosfere ordinatamente disposte nello spazio. In questo caso sono stati considerati diversi valori per il raggio e per la distanza tra le nanosfere. Il metodo *particle-in-cell* si è dimostrato essere lo strumento adeguato per lo studio del complesso sistema in esame. Le caratteristiche peculiari della fisica dell'interazione risultano essere ben catturate, come ad esempio l'assorbimento e la riflessione del laser da parte del bersaglio, la formazione della nuvola elettronica sul retro del bersaglio e i primi stadi del processo di accelerazione. Alcuni risultati numerici trovano qualitativamente riscontro nella letteratura sperimentale, come le particolari strutture ottenute nella distribuzione angolare degli ioni emessi e il tipico andamento esponenziale del loro spettro. Emerge chiaramente che la presenza della nanostruttura influisce significativamente sul processo di interazione, portando ad un incremento dell'assorbimento dell'energia del laser, il che a sua volta risulta in un incremento dell'energia massima degli ioni accelerati. Inoltre, le particolari proprietà della nanostruttura giocano un ruolo importante nel tentativo di migliorare il processo di accelerazione: schiume con parametri diversi portano a diverse prestazioni. In particolare, con le simulazioni 3D si è riscontrata una dipendenza dallo spessore medio della regione illuminata dall'impulso; con le simulazioni 2D si è riscontrata una dipendenza dalla dimensione e distanza delle nanosfere. In entrambi gli scenari sembra che, idealmente, una schiuma uniforme possa portare a migliori risultati a parità di tutti gli altri parametri. Essendo il setup delle simulazioni simile a quello di alcuni esperimenti in letteratura, è stato possibile effettuare confronti

tra risultati numerici e risultati sperimentali. Un accordo soddisfacente tra i due è stato osservato per esempio per quanto riguarda il fattore di guadagno nell'energia massima tra configurazioni convenzionale e avanzate. In generale, tuttavia, le simulazioni non possono essere predisposte in modo che i parametri del sistema simulato siano completamente realistici, neanche se il codice viene lanciato sui supercomputer di ultima generazione. I limiti nella scelta di questi parametri sono dovuti all'elevato costo computazionale che questo tipo di simulazioni comporta. Per questo è stato sviluppato un ulteriore livello di parallelismo all'interno del codice *piccante* sfruttando lo schema OpenMP mediante diverse strategie. Risultati soddisfacenti sono stati ottenuti per configurazioni di plasmi molto speciali con forti non uniformità, mentre promettenti risultati preliminari sono stati ottenuti su particolari architetture multicore.

Sia per il lavoro di simulazioni che per il lavoro di programmazione che sono riportati è stato utilizzato il codice particle-in-cell open-source *piccante*. Le simulazioni sono state effettuate sul cluster Intel Marconi (3D) e sul supercomputer Fermi (2D), mentre le performance del codice sono state testate sul cluster Intel Galileo, tutti e tre del centro di calcolo CINECA, Bologna.

Contents

Introduction	13
1 Elements of ion acceleration via superintense laser-plasma interaction	1
1.1 Introduction	1
1.2 High intensity laser	2
1.3 The plasma state	5
1.3.1 Fundamental plasma parameters	5
1.3.2 Kinetic model	7
1.4 Electromagnetic radiation-plasma interaction	14
1.4.1 Laser-particle interaction	15
1.5 Laser-driven ion acceleration	18
1.5.1 Acceleration mechanisms: target normal sheath acceleration	21
1.5.2 Nanostructured targets	22
2 Numerical methods for the Vlasov-Maxwell system	27
2.1 Phase space-grid based methods	27
2.1.1 Time splitting	29
2.1.2 Semi-Lagrangian methods	29
2.1.3 Spectral methods	31
2.2 The particle-in-cell method	32
2.2.1 Basic principles of the PIC scheme	32
2.2.2 Theoretical derivation	35
2.2.3 Problem definition and standard algorithms	36
2.2.4 Variations of the standard PIC method	46
2.3 Comparisons	47
2.4 Motivations and goals of this thesis work	48
3 Numerical simulations of laser-plasma interaction with nanostructured targets	51
3.1 Modelling nanostructured targets for laser-driven ion acceleration	51
3.2 Goals and methods of the numerical campaign	54

3.3	<i>piccante</i> : a particle-in-cell code	55
3.4	Results	57
3.4.1	Three-dimensional PIC simulations	57
3.4.2	Two-dimensional PIC simulations	71
3.4.3	Discussion	79
3.5	Computational issues	81
4	Code improvements	83
4.1	OpenMP parallelism	83
4.1.1	Parallelization strategies	84
4.1.2	Numerical tests	90
4.2	Charge-conserving current deposition method	96
4.2.1	Esirkepov algorithm	97
4.2.2	Numerical tests	100
4.3	Conclusions	104
5	Conclusions and perspectives	105
A	Klimontovich approach	109
	Bibliography	119

Introduction

Since the laser technology was invented in 1960, the interaction between a laser pulse and matter has been the focus of a great number of investigations. The regime of such interaction depends both on the laser pulse properties (duration, energy, focal spot, wavelength, intensity, coherence,...) and irradiated matter characteristics (material, density, thickness,...). Ongoing studies aim at controlling the interaction process by suitably tuning these parameters in a reproducible scheme.

The development of laser technology in the past decades led to higher and higher peak intensities, especially since the mid-eighties, when the amplification of low energy, short laser pulses was made possible thanks to the introduction of the *chirped pulse amplification* technique. The consequent availability of high intensity laser pulses opened the way to the exploration of new laser-matter interaction regimes, among which laser-driven particle acceleration takes place. Nowadays, state-of-the-art laser facilities can deliver pulses with high powers (hundreds of TW), ultra-short durations (tens of fs), focused down to small areas (few μm^2). These kinds of laser pulses, called *superintense*, are those employed in laser induced particle acceleration experiments.

In the year 2000, intense emission of protons with unique properties (multi-MeV energies, collimation,...) was observed for the first time in three separate experiments of laser-matter interaction with thin solid targets. Immediately the interest focused on the understanding of the underlying physics and on the possible technological applications. Indeed, accelerated ion bunches can be exploited in different contexts, all of which require an enhancement of the emitted beam properties observed thus far in laser-driven experiments, above all some of the spectra characteristics. A great deal of efforts have been and are still being made to optimize the acceleration process in order to obtain beam and target parameters of concrete interest for applications that require a reliable, efficient and robust laser-acceleration scheme. One option is to use special target configurations, such as double-layer targets that consist in a nanostructured, low density material, called *foam*, attached to a thin solid foil like those conventionally used in laser-driven ion acceleration experiments. Foam layers are added to foster the target absorption of laser energy, which ultimately leads to an enhancement of the acceleration performances. This can be accomplished only with

very low values of the foam (average) density, which has to be around few mg/cm^3 . Such low-density materials can be manufactured using the *pulsed laser deposition technique*, that allows to directly grow a foam layer on the solid foil. Because of the specific growth process inherent to the technique, low density layers result to be highly nanostructured, meaning that are constituted by aggregates of nanoparticles (ensembles of particles of size about tens of nm) that assemble in particular structures on the mesoscale (scale lengths $\gtrsim \mu\text{m}$).

Many diverse, complex physical mechanisms come at play during the interaction between a superintense laser pulse and matter. In any case, the irradiated matter strongly ionizes and rapidly turns to plasma, meaning that can be considered as a collection of charged particles in electromagnetic interaction with each other. Temperatures, densities, time scales and length scales of interest for laser-driven ion acceleration are such that the behaviour of the plasma is predominantly collective because long-range electromagnetic interactions prevail over those on the microscopic scales. Theoretical investigations of such systems mainly rely on the kinetic plasma theory. In particular, the most suitable mathematical model for their description is given by the relativistic Vlasov equation coupled with Maxwell system. Reductions and simplifications of this model to allow analytical analysis are possible, but not adequate to perform in-depth studies of real configurations, such as those under investigation in the framework of laser-based ion acceleration experiments. In this case, numerical tools are required. Among them, the so-called *particle-in-cell* method is particularly suitable for the kinetic simulation of laser-plasma interactions.

The scope of the present thesis is the numerical investigation of ion acceleration by means of the interaction between superintense, ultra-short laser pulses and nanostructured plasmas. This work has the general goal to explore the feasibility of using particle-in-cell codes to gain insights on the physics of the interaction through expensive, large scale three and two-dimensional particle-in-cell simulations representative to some degree of the experiments, paying particular attention to the modelling of the target morphology and structure together with the choice of the foam properties. Moreover, the will of performing larger sets of simulations with parameters closer to the experimental ones led to the necessity of improving the code by developing further optimizations.

The code used to perform all simulations is the open source, massively parallel, particle-in-cell code *piccante*. Simulations were performed on the Intel cluster Marconi and Fermi supercomputer, while code performances have been tested on the Intel cluster Galileo, all of them hosted at CINECA, Bologna, Italy.

The thesis is organized as follows.

- Chapter 1 – *Elements of ion acceleration via superintense laser-plasma interaction* – frames the problem from the physical and mathematical point of view. After a

brief presentation of high intensity lasers, the concept of plasma is introduced along with the adopted mathematical model for its description, i.e. the relativistic Vlasov-Maxwell system. Following, some general features of the interaction between a laser pulse and a plasma are discussed. Lastly, the basic elements specific of laser based ion acceleration are presented, with particular attention towards novel target concepts consisting in nanostructured materials deposited on solid foils.

- Chapter 2 – *Numerical methods for the Vlasov-Maxwell system* – presents the two main classes of methods used for the approximated solution of the Vlasov-Maxwell system: grid-based methods and particle methods. A short overview of the former is followed by an in-depth description of the particle-in-cell method, the most exploited among particle methods and the one employed in this work. The main differences between the two approaches are highlighted. Motivations and goals of the present thesis are detailed at the end of this chapter.
- Chapter 3 – *Numerical simulations of laser-plasma interaction with nanostructured targets* – provides an extended discussion on the performed simulations. The adopted models for nanostructured targets description are presented. Then, the aims of the numerical campaigns and the tools used to run the simulations are specified. The simulation results, both three-dimensional and two-dimensional, are presented with many graphical supports. Lastly, the main issues concerning the computational resources are addressed also to point out the necessity of code improvements.
- Chapter 4 – *Code improvements* – is devoted to the development of two new features in the PIC code *piccante*. First, an additional parallelization strategy, based on the memory-shared paradigm OpenMP. Second, a new charge-conserving current deposition scheme, i.e. the Esirkepov method.
- Chapter 5 – *Conclusions and future perspectives* – discusses the over-all results of the work and highlights possible further developments along with ideas for potential future studies.

Chapter 1

Elements of ion acceleration via superintense laser-plasma interaction

Superintense laser induced ion acceleration is an exciting research field both for the applicative potential and for the investigation of particular regimes of radiation-matter interaction. Nowadays, efforts focus on the enhancement of the acceleration process by employing different strategies, such as the exploitation of special targets. This chapter presents the basic concepts needed to frame such problem. After a brief introduction, some issues about lasers, plasmas and their interaction are discussed in sections 1.2, 1.3 and 1.4 respectively. Lastly, section 1.5 concerns those aspect of laser-based ion acceleration relevant for this thesis work.

1.1 Introduction

Since the invention of the laser in 1960, a great interest has been addressed towards the study of laser-matter interaction for a variety of scientific and technological applications. Here the focus is mainly on laser-matter interaction at very high field intensities. Nowadays, state-of-the-art laser facilities can produce laser pulses delivering extremely high powers (up to ~ 1 PW) on very small areas (of size ~ 10 - $100 \mu\text{m}$) for very short times (tens fs - ~ 1 ps depending on the laser technology). These so-called *ultra-intense* and *ultra-short* laser pulses, when impinging onto a target, rapidly (i.e. in a few laser cycles) ionize matter, producing a plasma. Thus laser-matter interaction in these regimes means invariably laser interaction with a plasma.

The interaction between high intensity laser pulses and plasmas can be exploited for a

tremendous number of applications, one of these being laser-based particle acceleration. Indeed, a unique feature of plasmas is that they are able to sustain huge electric fields, in the TV/m range, that can be with suitable strategies exploited for charged particle acceleration. The principle for accelerating charged particles using plasmas was first proposed by Veksler in 1957 [1]. It is called *coherent acceleration* because the magnitude of the accelerating field acting on each particle is proportional to the number of particles being accelerated. It is worth mentioning that due to the great difference between electron and proton masses, the schemes for electron and ion acceleration are based on different concepts.

Typically, in electron acceleration schemes the laser pulse is made to interact with a gas jet, producing a low density plasma in which the pulse can propagate for long distances. On the contrary laser-driven ion acceleration schemes foresee the irradiation of a thin solid foil, which becomes a solid density plasma (for lasers with a temporal duration of few 10s fs there is no time for hydrodynamic expansion). This fact relies on basic concepts of radiation-plasma interaction that will be presented later in the chapter.

Thus far it has been possible to generate electron beams with energies up to 4.25 GeV on few centimetres distances [2] and proton bunches with energies up to 85 MeV on few microns distances [3]. Such beams have exhibited interesting properties, albeit they need to be improved for future applications. Many ongoing researches aim at enhancing these properties, e.g. trying to reach higher kinetic energies of the accelerated particles and a higher number of accelerated particles, with relaxed experimental configurations together with a higher degree of controllability and reproducibility. As will be explained in the following, in the framework of ion acceleration one way towards this goal is represented by the exploitation of special multi-layer targets. In general, from the applicative point of view, the compactness of the scheme makes laser-induced particle acceleration an attractive and affordable alternative to larger, more expensive conventional accelerators [4]. Moreover, the interest in such phenomena also comes from the opportunity of performing fundamental physics investigations.

1.2 High intensity laser

A laser (light amplification by stimulated emission of radiation) is a device that acts as a source of coherent electromagnetic radiation, localized both in space and time. Coherence is a property of the spatial and temporal behaviour of the phase of the emitted radiation that is related to some of its peculiar features. Spatially, it allows the emitted light to be focused to a small (down to the diffraction limit) spot and collimated over great distances; temporally, it allows the emission in a very narrow spectrum, so that the wave is approximately monochromatic. The electric field corresponding to a laser pulse can have

very high amplitude with important spatial non-uniformities, which give rise to many complex nonlinear effects. The main characteristics of a laser pulse are duration, energy, spot size and wavelength (indicated with τ , \mathcal{E} , σ and λ respectively). The energy of the pulse can be estimated as that contained in the corresponding electromagnetic field, using the equipartition between electric and magnetic components:

$$\mathcal{E} \approx \frac{|\mathbf{E}|^2}{4\pi} \sigma c \tau,$$

where $\sigma c \tau$ is the volume occupied by the pulse. An important quantity is the *intensity* of the pulse, i.e. the power flowing through the unit surface:

$$I = \frac{\mathcal{E}}{\sigma \tau} \approx \frac{|\mathbf{E}|^2}{4\pi} c.$$

Many laser systems produce pulses which can be approximated as having a Gaussian intensity envelope, modelled as:

$$I(x, y, z, t) = I_0 \left(e^{-2(x^2+y^2)/w_0^2} e^{-(z-ct)^2/\sigma_t^2} \right),$$

where z is the propagation direction, the *waist* w_0 is the size of the focal spot and σ_t is related to the pulse duration τ , defined as the full-width-half-maximum of the intensity profile, $\tau^2 = 4 \ln 2 \sigma_t^2$.

Many laser technologies exist which produce pulses with different characteristics (see Svelto [5]). As far as high power laser systems are of concern, three main technologies exist: Titanium:Sapphire lasers, CO₂ lasers, Nd:YAG lasers. Pulses with durations down to tens of fs can be provided by Ti:Sapphire lasers, whereas durations in the ps and hundreds of fs ranges can be obtained with CO₂ and Nd:YAG lasers respectively. The highest peak powers reached with Ti:Sapphire and Nd:YAG lasers are about 1 PW, while less than 100 TW with CO₂ lasers. The pulse has a wavelength of approximately 0.8 μm for Ti:Sapphire systems, 1 μm for Nd:YAG and 10 μm for CO₂. Typically, in laser-driven ion acceleration experiments present-day solid state lasers based on the Titanium:Sapphire technology are employed, which can generate pulses with durations about tens of fs carrying tens of J, focused to a few wavelengths spot sizes. This implies powers in excess of 1 PW and intensities up to 10^{22} W/cm². These ultra-intense laser systems are based on the so-called *Chirped Pulse Amplification* (CPA), which was introduced in the mid-eighties, fostering the race towards higher and higher intensities (see [6]). Briefly, CPA is a technique used to amplify a short, low-energy laser pulse (see figure 1.1). In a first step, the duration of the laser pulse is increased by dispersing in time its spectral components. The resulting pulse has lower power and it is easier and safer to amplify. Then, the pulse undergoes different amplification stages that increment its energy by several orders of magnitude (from nJ up to tens of J). Lastly, the pulse is compressed again in time so that its final power can reach hundreds of TW. As of today, ultra-intense ($> 10^{21}$ W/cm²) pulses can only be

obtained using the CPA scheme and the world record for the highest focused intensity is 2×10^{22} W/cm², reached by the HERCULES Petawatt Laser at the University of Michigan. Currently, several laser facilities worldwide are able to deliver pulses with powers between 200 TW and 1 PW [7], while efforts are being made to overcome the PW limit. The European project ELI (Extreme Light Infrastructure), the Shanghai Institute of Optics and Fine Mechanics (SIOM) and the Laboratory for the Use of Intense Lasers at École Polytechnique (Paris-Saclay) are working on pioneering 10 PW laser systems.




Figure 1.1: Principle of chirped pulse amplification.

Ultra-intense laser pulses carry intensities well above the threshold needed to ionize matter. Indeed, matter can be easily ionized, hence becoming a plasma, through the interaction with an intense enough laser pulse. The order of magnitude of the required intensity can be estimated with a simple computation. For instance, the characteristic Coulombian field of an hydrogen atom is about e/r^2 , where e is the elementary charge in statcoulombs and r is the Bohr radius; the corresponding intensity is 10^{14} W/cm². With intensities greater than 10^{16} W/cm², the electric field of the pulse is able to suppress the Coulomb barrier of any atom.

For even higher intensity values, relativistic effects become non negligible primarily on electrons dynamic. To understand whether the electron motion is relativistic or not the normalized vector potential a_0 is used and is defined as:

$$a_0 = \frac{eA_0}{m_e c^2} = \sqrt{\frac{e^2 \lambda^2 I_0}{2\pi m_e c^3}} = 0.85 \lambda [\mu\text{m}] \sqrt{I_0 [10^{18} \text{ W/cm}^{-2}]}, \quad (1.1)$$

where A_0 is the peak potential vector and I_0 is the peak intensity of the laser pulse. Relativistic effects are not negligible if $a_0 \simeq 1$, namely when an electron acquires a kinetic energy comparable with its rest energy ($m_e c^2 = 0.511$ MeV) in a single laser cycle. In terms of the so-called laser *irradiance* this means that $I_0 \lambda^2 \simeq 10^{18} \frac{\text{W} \mu\text{m}^2}{\text{cm}^2}$. For ultra-intense lasers

this means that at intensities greater than 10^{18} W/cm² the electrons dynamic is relativistic. Even higher intensities, greater than 10^{24} W/cm², would make the ions dynamic relativistic as well; anyhow such intensities are well beyond present-day capabilities.

Summarizing, when an ultra-intense laser pulse interacts with matter it produces a plasma whose dynamic is relativistic, at least for what concerns the electrons. Besides, many non-linear effects arise from the laser-plasma coupling, making the underlying physics extremely complex. Lastly, it is worth mentioning that laser-based particle acceleration requires relativistic intensities, i.e. $I \geq 10^{18}$ W/cm². Such intensities have been reached for the first time in the nineties; it was then when the thrill for this research field begun to arise.

1.3 The plasma state

As a first, simplified, definition, a *plasma* is a globally neutral medium consisting of a collection of free charged particles that organize themselves interacting with each other through the *self-consistent* electromagnetic field that arises as a consequence of their motion. Tonks and Langmuir first introduced the word "plasma" in 1929 [8] to designate the high density, neutral region of an arc discharge in a ionized gas. A strongly ionized gas, where most of the atoms are broken down into electrons and ions, is a typical example of a plasma. Many different plasmas exist in the universe. The 99.9% of the visible matter in the universe is in the state of plasma: stellar cores and surfaces, solar wind, interstellar space, ionosphere. Also laboratory plasmas exist, such as tokamak plasmas or laser induced plasmas.

Different *populations* or *species* of charged particles – and possibly neutrals – comprise a plasma, being a large fraction of the electrons of the system separated from their nuclei. The species are free to move and consequently generate electromagnetic fields through their charge and current densities. It is a many-body system of strongly interacting entities whose dynamics is dominated by the collective behaviour through the self-consistent electromagnetic forces. A self-consistent description is required, meaning that electromagnetic field dynamics and species dynamics are coupled. Therefore, a plasma may be considered as the ensemble of self-consistent electromagnetic field and charged populations.

1.3.1 Fundamental plasma parameters

Every plasma is characterized by a set of parameters that may differ by many orders of magnitude from one plasma to another. This is a unique feature: while the characteristic parameters vary in extremely large ranges, the underlying physical laws are the very same. Characteristic scale lengths can range from nanometres up to metres, densities from 1 cm^{-3} to 10^{30} cm^{-3} and temperatures from 10^4 K ("cold plasmas") to GK (see table 1.2).

The most relevant parameters for plasma description are related to its collective behaviour. At least two different types of energy contributions can always be associated to a plasma: kinetic and electromagnetic. During the evolution of the system these terms are in a continuous competition. This occurs because of the spontaneous charge separation that can be observed in a plasma at a given temperature. The characteristic length corresponding to such charge displacement is known as *Debye length* and is defined by:

$$\lambda_D = \left(4\pi e^2 \sum_{a=1}^N \frac{n_{a,0} Z_a^2}{T_a} \right)^{-1/2}, \quad (1.2)$$

where $n_{a,0}$ is the density in the unperturbed configuration, T_a is the temperature and Z_a the atomic number of the a -th species. A *Debye sphere* is a sphere of radius λ_D centred on any particle, inside of which the short-range Coulombian interactions, i.e. collisions, are relevant. On the other hand, every particle is shielded from the others on distances greater than λ_D : on scales $> \lambda_D$ only long-range electromagnetic interactions are significant. Each particle can be thought as *dressed* with its own Debye sphere, forming a new entity that is called *quasi-particle* because it behaves like an ordinary particle for what concerns some of its properties. Therefore the whole plasma can be thought as made of a collection of overlapping quasi-particles that are neutral and weakly interacting. A Debye length can be defined for each species, for example for electrons:

$$\lambda_{De} = \sqrt{\frac{T_e}{4\pi n_{0,e} e^2}}. \quad (1.3)$$

Using the thermal velocity

$$v_{te} = \sqrt{\frac{T_e}{m_e}}$$

it is possible to define the electron *plasma frequency*:

$$\omega_{pe} = \frac{v_{te}}{\lambda_{De}} = \sqrt{\frac{4\pi n_{0,e} e^2}{m_e}},$$

which is the characteristic frequency of the electron dynamics, related to an harmonic-like process. Among all populations, the electron is the one related to the shortest times scales, since the plasma frequency depends on the $-1/2$ power of the mass. Another important parameter is the so called *plasma parameter*, defined as the inverse number of particles per Debye length:

$$g = \frac{1}{N_D} = \frac{1}{\frac{4}{3}\pi\lambda_D^3 n} \sim \left(\frac{3T/2}{e^2/n^{-1/3}} \right)^{3/2},$$

which, roughly, is the $3/2$ power of the ratio between a particle thermal energy and the potential energy related to the Coulomb interaction between two particles inside a Debye sphere. g is a measure of the degree of non-ideality of the plasma: the more g is close to zero, the more the correlations on small scale lengths ($< \lambda_D$) are negligible. At the limit $g = 0$

the plasma is ideal, i.e. a perfect gas characterized by only thermal energy, subjected to the self-consistent electromagnetic field. If $g \simeq 0$, then the plasma is sufficiently dilute and hot; moreover, collisional effects may be considered through a perturbative approach.

Plasma Type	$n \text{ cm}^{-3}$	$T \text{ eV}$	$\omega_{pe} \text{ sec}^{-1}$	$\lambda_D \text{ cm}$	$n\lambda_D^3$	$\nu_{ei} \text{ sec}^{-1}$
Interstellar gas	1	1	6×10^4	7×10^2	4×10^8	7×10^{-5}
Gaseous nebula	10^3	1	2×10^6	20	8×10^6	6×10^{-2}
Solar Corona	10^9	10^2	2×10^9	2×10^{-1}	8×10^6	60
Diffuse hot plasma	10^{12}	10^2	6×10^{10}	7×10^{-3}	4×10^5	40
Solar atmosphere, gas discharge	10^{14}	1	6×10^{11}	7×10^{-5}	40	2×10^9
Warm plasma	10^{14}	10	6×10^{11}	2×10^{-4}	8×10^2	10^7
Hot plasma	10^{14}	10^2	6×10^{11}	7×10^{-4}	4×10^4	4×10^6
Thermonuclear plasma	10^{15}	10^4	2×10^{12}	2×10^{-3}	8×10^6	5×10^4
Theta pinch	10^{16}	10^2	6×10^{12}	7×10^{-5}	4×10^3	3×10^8
Dense hot plasma	10^{18}	10^2	6×10^{13}	7×10^{-6}	4×10^2	2×10^{10}
Laser Plasma	10^{20}	10^2	6×10^{14}	7×10^{-7}	40	2×10^{12}

Figure 1.2: Approximate magnitudes in some typical plasmas. Reprinted from the NRL Plasma Formulary. Here "laser plasmas" refers to the interaction with long pulses (ns durations) and relatively small intensities (10^{14} W/cm^2).

1.3.2 Kinetic model

A plasma can be described using several mathematical models, depending on the problem under consideration and on the required level of detail. The most complete, solvable description is inherited from the kinetic theory (the actual most complete theory is not solvable even with state-of-the-art supercomputers). The system under examination consists of a macroscopically large number of interacting particles, therefore a statistical approach can be adopted to build a macroscopic description of the system from the microscopic behaviour of the single entities. To derive the kinetic model for a plasma the kinetic theory must be adapted to the self-consistent electromagnetic interaction. One option is to follow the approach first proposed by Klimontovich. Other arguments are possible; among them, Klimontovich's is particularly enlightening in view of the numerical schemes used to solve the plasma dynamic presented in chapter 2.

Klimontovich approach

Consider a plasma made of N species each with mass m_a , charge q_a and N_a particles. Under the non-relativistic approximation, the equation of motion of the i -th particle of kind a is given by:

$$\begin{cases} \frac{d\mathbf{r}_{i,a}}{dt} = \mathbf{v}_{i,a} \\ m_a \frac{d\mathbf{v}_{i,a}}{dt} = q_a \left(\mathbf{E}_{micr} + \frac{\mathbf{v}_{i,a}}{c} \times \mathbf{B}_{micr} \right), \end{cases} \quad (1.4)$$

where $\mathbf{r}_{i,a}(t)$, $\mathbf{v}_{i,a}(t)$ are the particle position and velocity respectively, while $\mathbf{E}_{micr} = \mathbf{E}_{micr}(\mathbf{r}, t)$ and $\mathbf{B}_{micr} = \mathbf{B}_{micr}(\mathbf{r}, t)$ are the microscopic electric and magnetic fields expressed in Gaussian units, solutions of the Maxwell equations:

$$\begin{cases} \text{curl } \mathbf{E}_{micr} = -\frac{1}{c} \frac{\partial \mathbf{B}_{micr}}{\partial t} \\ \text{curl } \mathbf{B}_{micr} = \frac{4\pi}{c} \mathbf{J}_{micr} + \frac{1}{c} \frac{\partial \mathbf{E}_{micr}}{\partial t} \\ \text{div } \mathbf{E}_{micr} = 4\pi \rho_{micr} \\ \text{div } \mathbf{B}_{micr} = 0, \end{cases} \quad (1.5)$$

where the sources are given by the charge and current densities of the external charges together with the plasma charged populations:

$$\begin{aligned} \rho_{micr}(\mathbf{r}, t) &= \rho_{ext} + \sum_{a=1}^N q_a \sum_{i=1}^{N_a} \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) \\ \mathbf{J}_{micr}(\mathbf{r}, t) &= \mathbf{J}_{ext} + \sum_{a=1}^N q_a \sum_{i=1}^{N_a} \mathbf{v}_{i,a} \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)). \end{aligned}$$

Plasma dynamics is completely described by the solution of (1.4) and (1.5). This is a system of a huge number ($\sum_{a=1}^N N_a \gg 1$) of extremely complex equations, therefore it is impossible to make direct use of this model. Before deriving a simpler description it is useful to write an equivalent formulation of the problem. To do so, introduce the *phase space* of coordinates (\mathbf{r}, \mathbf{v}) , position and velocity Eulerian variables respectively. As a remark, the velocity variable could be replaced with the momenta variable $\mathbf{p}_a = m_a \mathbf{v}$. In some cases it is especially convenient, as when relativistic effects are taken into account. From time to time, the most suitable option will be adopted; for the sake of simplicity here the \mathbf{v} variable is used. Now, on the phase space (\mathbf{r}, \mathbf{v}) define the *microscopic distribution*

function for the a -th plasma component as

$$f_{micr,a}(\mathbf{r}, \mathbf{v}, t) = \sum_{i=1}^{N_a} \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) \delta(\mathbf{v} - \mathbf{v}_{i,a}(t))$$

so that

$$f_{micr,a}(\mathbf{r}, \mathbf{v}, t) d\mathbf{r} d\mathbf{v}$$

is the number of particles of kind a that are located in a neighbourhood of the point (\mathbf{r}, \mathbf{p}) of the phase space at time t . The microscopic state of the system is determined by these functions, which become the unknowns to an equivalent formulation of the original problem. Expressing the sources in Maxwell equations in terms of $f_{micr,a}$ as

$$\begin{aligned} \rho_{micr}(\mathbf{r}, t) &= \rho_{ext} + \sum_{a=1}^{N_a} q_a \int f_{micr,a}(\mathbf{r}, \mathbf{v}, t) d\mathbf{v} \\ \mathbf{J}_{micr}(\mathbf{r}, t) &= \mathbf{J}_{ext} + \sum_{a=1}^{N_a} q_a \int \mathbf{v} f_{micr,a}(\mathbf{r}, \mathbf{v}, t) d\mathbf{v}, \end{aligned}$$

and using the equations of motion (1.4) it is possible to derive the *Klimontovich equation*, i.e. the continuity equation for the evolution of $f_{micr,a}$ in the phase space:

$$\frac{\partial f_{micr,a}}{\partial t} + \operatorname{div}_{\mathbf{r}}(\mathbf{v} f_{micr,a}) + \operatorname{div}_{\mathbf{v}} \left[\frac{q_a}{m_a} \left(\mathbf{E}_{micr} + \frac{\mathbf{v}}{c} \times \mathbf{B}_{micr} \right) f_{micr,a} \right] = 0. \quad (1.6)$$

For a simple but complete derivation of Klimontovich equation refer to appendix A. This is an equivalent formulation of the initial problem, therefore still unsolvable, but it can be simplified using a suitable average operation. Consider an ensemble average operation $\langle \cdot \rangle$ that is an average over all possible microscopic configurations that correspond to the same macroscopic state. This allows the decomposition of a microscopic quantity in two contributes: one related to the average behaviour, the other to the fluctuating behaviour. Such operation, besides its statistical relevance, is able to filter the spatial scales so that only those characteristic of the problem remain at play. Let the *distribution function* f_a be the ensemble average of the microscopic distribution function and in the same fashion define the mean electromagnetic field \mathbf{E} , \mathbf{B} :

$$\begin{aligned} f_{micr,a} &= \langle f_{micr,a} \rangle + \delta f_a = f_a + \delta f_a, \\ \mathbf{E}_{micr} &= \langle \mathbf{E}_{micr} \rangle + \delta \mathbf{E} = \mathbf{E} + \delta \mathbf{E}, \\ \mathbf{B}_{micr} &= \langle \mathbf{B}_{micr} \rangle + \delta \mathbf{B} = \mathbf{B} + \delta \mathbf{B}. \end{aligned}$$

Applying the average operation to the Maxwell equations yields:

$$\left\{ \begin{array}{l} \text{curl } \mathbf{E} = -\frac{1}{c} \frac{\partial \mathbf{B}}{\partial t} \\ \text{curl } \mathbf{B} = \frac{4\pi}{c} \mathbf{J} + \frac{1}{c} \frac{\partial \mathbf{E}}{\partial t} \\ \text{div } \mathbf{E} = 4\pi \rho \\ \text{div } \mathbf{B} = 0, \end{array} \right. \quad (1.7)$$

where

$$\rho(\mathbf{r}, t) = \rho_{ext} + \sum_{a=1}^N q_a \int f_a(\mathbf{r}, \mathbf{v}, t) d\mathbf{v}, \quad (1.8)$$

$$\mathbf{J}(\mathbf{r}, t) = \mathbf{J}_{ext} + \sum_{a=1}^N q_a \int \mathbf{v} f_a(\mathbf{r}, \mathbf{v}, t) d\mathbf{v}. \quad (1.9)$$

Applying the same average operation to the Klimontovich equation yields to the *Boltzmann equation* for a plasma:

$$\frac{\partial f_a}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{r}} f_a + \frac{q_a}{m_a} \left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B} \right) \cdot \nabla_{\mathbf{v}} f_a = -\frac{q_a}{m_a} \left\langle \left(\delta \mathbf{E} + \frac{\mathbf{v}}{c} \times \delta \mathbf{B} \right) \cdot \nabla_{\mathbf{v}} \delta f_a \right\rangle \quad (1.10)$$

It is possible to recognize two different types of electromagnetic interaction between particles. One is the macroscopic, averaged, long-range interaction associated to the collective motion of relatively large volumes of particles through the mean self-consistent field (left-hand-side of (1.10)). The other is the microscopic, chaotic, short-range interaction, related to the collisions between particles (right-hand-side of (1.10)). Having applied an average operation, collisional effects manifest through a correlation term, $C_a = -\frac{q_a}{m_a} \left\langle \left(\delta \mathbf{E} + \frac{\mathbf{v}}{c} \times \delta \mathbf{B} \right) \cdot \nabla_{\mathbf{v}} \delta f_a \right\rangle$, called *total collision integral*. Being C_a unknown, the problem is not close. For a strongly ionized plasma at high temperature the easiest and strongest closure consists in completely neglecting collisional effects, which means considering uncorrelated fluctuations of the distribution function and electromagnetic field (in other words the plasma is ideal and $g = 0$). This approximation is meaningful as long as the evolution of the system is of concern for short times with respect to the relaxation time, i.e. the time scale over which the perturbed system returns to the equilibrium configuration. This procedure leads to the *non-relativistic Vlasov equation* for the distribution function of a collisionless plasma:

$$\frac{\partial f_a}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{r}} f_a + \frac{q_a}{m_a} \left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B} \right) \cdot \nabla_{\mathbf{v}} f_a = 0. \quad (1.11)$$

It is worth noting that Klimontovich equation and Vlasov equation are formally identical from the mathematical point of view. Equations (1.7), coupled with equation (1.11)

comprise the mathematical model that describes the kinetic, non-relativistic dynamics of a collisionless plasma. The unknowns of the problem are the distribution function for each species f_a and the electromagnetic field \mathbf{E} , \mathbf{B} . Of course, for the sake of well-posedness, they need to be equipped with appropriate boundary and initial conditions. Overall, the *Vlasov-Maxwell system* reads:

$$\left\{ \begin{array}{l} \frac{\partial f_a}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{r}} f_a + \frac{q_a}{m_a} \left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B} \right) \cdot \nabla_{\mathbf{v}} f_a = 0 \\ \text{curl } \mathbf{E} = -\frac{1}{c} \frac{\partial \mathbf{B}}{\partial t} \\ \text{curl } \mathbf{B} = \frac{4\pi}{c} \mathbf{J} + \frac{1}{c} \frac{\partial \mathbf{E}}{\partial t} \\ \text{div } \mathbf{E} = 4\pi\rho \\ \text{div } \mathbf{B} = 0, \end{array} \right. \quad (1.12)$$

where the sources ρ and \mathbf{J} are given by (1.8) and (1.9).

Conservation properties of the Vlasov-Maxwell system

From the Vlasov-Maxwell system, many important conservation properties can be deduced. For the sake of simplicity only one species is considered. The solution of Vlasov equation can be expressed using the *characteristics*, i.e. the trajectories of the particles in phase space, solutions of the following differential system:

$$\left\{ \begin{array}{l} \frac{d\mathbf{R}(t)}{dt} = \mathbf{V}(t) \\ \frac{d\mathbf{V}(t)}{dt} = \frac{q}{m} \left(\mathbf{E}(\mathbf{R}(t), t) + \frac{\mathbf{V}(t)}{c} \times \mathbf{B}(\mathbf{R}(t), t) \right). \end{array} \right. \quad (1.13)$$

Because the distribution function satisfies Vlasov equation, it is conserved along the characteristics. Hence, Vlasov equation can be expressed in the form of a conservation law as follows:

$$\frac{d}{dt} f(\mathbf{R}(t), \mathbf{V}(t), t) = 0.$$

As time passes by, the characteristics may become close to each other, meaning that also the regions where f takes different values become close to each other. This mechanism can continue on finer and finer scale lengths. Physically this is a consequence of the fact that Vlasov-equation is non-dissipative and entropy-conserving. The action of the

Lorentz force does not vanish in time, rather it engages smaller scale lengths in an energy cascade process. This fact is known as *phase space filamentation* because of the filamentous structures observed in the phase space, (see figure 1.3).




Figure 1.3: Example of phase space filamentation. Reprinted from [9].

Let $(\mathbf{R}(t; \mathbf{r}, \mathbf{v}, s), \mathbf{V}(t; \mathbf{r}, \mathbf{v}, s))$ be the unique solution of (1.13) at time t that takes the value (\mathbf{r}, \mathbf{v}) at time s , then

$$f(\mathbf{r}, \mathbf{v}, t) = f_0(\mathbf{R}(0; \mathbf{r}, \mathbf{v}, t), \mathbf{V}(0; \mathbf{r}, \mathbf{v}, t)),$$

where f_0 is the initial distribution function. This means that f can be reconstructed from its initial values. This property implies the following maximum principle:

$$0 \leq f(\mathbf{r}, \mathbf{v}, t) \leq \max_{(\mathbf{r}, \mathbf{v})} f_0(\mathbf{r}, \mathbf{v}).$$

Integrating Vlasov equation on the whole phase space it is obvious that the total number of particles is conserved during the evolution of the system:

$$\frac{d}{dt} \int f(\mathbf{r}, \mathbf{v}, t) d\mathbf{r} d\mathbf{v} = 0 \quad \Longleftrightarrow \quad \int f(\mathbf{r}, \mathbf{v}, t) d\mathbf{r} d\mathbf{v} = N.$$

More in general, multiplying Vlasov equation by f^{p-1} and integrating in the whole phase space the conservation of all L^p norms of f can be deduced:

$$\frac{d}{dt} \int f(\mathbf{r}, \mathbf{v}, t)^p d\mathbf{r} d\mathbf{v} = 0.$$

Lastly, the equation can be written in conservative form as

$$\frac{\partial f}{\partial t} + \text{div}_{\mathbf{r}, \mathbf{v}}(\mathbf{F}f) = 0, \quad \mathbf{F} = \left(\mathbf{v}, \mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B} \right)^T, \quad \text{div}_{\mathbf{r}, \mathbf{v}} \mathbf{F} = 0.$$

Relativistic description

In many situations relativistic effects cannot be neglected because particles energies $E = m\gamma c^2$ may be comparable to the rest energy. This is often true for the electron populations of plasmas interacting with high intensity laser pulses as explained in section 1.2. Since

Maxwell equations are relativistically correct, only the description of the particles dynamics must change. The equations of motion become:

$$\begin{cases} \frac{d\mathbf{r}_{i,a}}{dt} = \frac{\mathbf{P}_{i,a}}{m_a\gamma_{i,a}} \\ \frac{d\mathbf{P}_{i,a}}{dt} = q_a \left(\mathbf{E}_{micr} + \frac{\mathbf{P}_{i,a}}{cm_a\gamma_{i,a}} \times \mathbf{B}_{micr} \right), \end{cases} \quad (1.14)$$

where

$$\begin{aligned} \mathbf{P}_{i,a} &= m_a\gamma_{i,a}\mathbf{v}_{i,a} \\ \gamma_{i,a} &= \sqrt{1 + \frac{\mathbf{P}_{i,a} \cdot \mathbf{P}_{i,a}}{m_a^2c^2}}. \end{aligned}$$

A balance equation for the energy of the particle $\gamma_{i,a}m_ac^2$ can be obtained by multiplying the momentum equation by the particle velocity:

$$\mathbf{v}_{i,a}(t) \cdot \frac{d\mathbf{P}_{i,a}(t)}{dt} = m_ac^2 \frac{d\gamma_{i,a}}{dt} = q_a\mathbf{v}_{i,a}(t) \cdot \mathbf{E}(\mathbf{r}_{i,a}, t). \quad (1.15)$$

Similarly to the non-relativistic case, consider the phase space of Eulerian coordinates $(\mathbf{r}, \mathbf{p}_a)$, position and a -th species momenta variables respectively, and introduce the distribution function $f_a = f_a(\mathbf{r}, \mathbf{p}_a, t)$. Using the Klimontovich approach, the relativistic Vlasov equation can be derived (see A for a simple derivation or Groot [10] for deeper insights):

$$\frac{\partial f_a}{\partial t} + \frac{\mathbf{p}_a}{\gamma_a m_a} \cdot \nabla_{\mathbf{r}} f_a + q_a \left(\mathbf{E} + \frac{\mathbf{p}_a}{m_a \gamma_a c} \times \mathbf{B} \right) \cdot \nabla_{\mathbf{p}_a} f_a = 0, \quad (1.16)$$

where \mathbf{p}_a and γ are the linear momentum of the a -th population and the relativistic factor defined by:

$$\begin{aligned} \mathbf{p}_a &= m_a\gamma_a\mathbf{v} \\ \gamma_a &= \frac{1}{\sqrt{1 - p_a^2/m_a^2c^2}}. \end{aligned}$$

Hence, the relativistic Vlasov-Maxwell system is:

$$\begin{cases} \frac{\partial f_a}{\partial t} + \frac{\mathbf{p}_a}{\gamma_a m_a} \cdot \nabla_{\mathbf{r}} f_a + q_a \left(\mathbf{E} + \frac{\mathbf{p}_a}{m_a \gamma_a c} \times \mathbf{B} \right) \cdot \nabla_{\mathbf{p}_a} f_a = 0 \\ \text{curl } \mathbf{E} = -\frac{1}{c} \frac{\partial \mathbf{B}}{\partial t} \\ \text{curl } \mathbf{B} = \frac{4\pi}{c} \mathbf{J} + \frac{1}{c} \frac{\partial \mathbf{E}}{\partial t} \\ \text{div } \mathbf{E} = 4\pi\rho \\ \text{div } \mathbf{B} = 0. \end{cases} \quad (1.17)$$

Both non-relativistic and relativistic Vlasov-Maxwell systems are too complex to be investigated with analytical theories without introducing strong approximations. The only option to examine the behaviour of real plasmas is to rely on numerical tools. The numerical approximation of the problem is the focus of chapter 2.

As a final remark, in some scenarios (e.g. extremely dense plasmas) not only relativistic effects, but also quantum effects need to be taken into account. This represents a non-trivial issue that goes far beyond the scope of the present work.

1.4 Electromagnetic radiation-plasma interaction

A fundamental issue in the physics of radiation-matter interaction is whether an electromagnetic wave, for instance a laser pulse, can propagate through a given medium or not. Plasmas are dispersive mediums characterized by several length and time scales, therefore their response to external perturbations may be delayed and non-local; furthermore it strongly depends on the relations between perturbation (e.g. external fields) properties and plasma parameters. In general an electromagnetic wave propagating within a plasma must satisfy Maxwell laws together with Kramers-Kronig relations, i.e. the dispersion relation (see Jackson [11]). Under suitable hypotheses, these equations lead to a simple condition between the wave frequency and wave vector. In particular, consider a non-magnetized, unperturbed, collisionless, non-relativistic, spatially non-dispersive plasma. Assuming linear response to small perturbations, i.e. linear theory for Maxwell equations in a material medium, a plane monochromatic electromagnetic wave with frequency ω and wave vector \mathbf{k} can be sustained by the plasma only if it is a transverse wave and:

$$\omega^2 = \omega_p^2 + |\mathbf{k}|^2 c^2,$$

or if it is a longitudinal wave and:

$$\omega = \omega_p.$$

As a remark, no spatial dispersion means that the spatial variations of the wave are greater than the characteristic lengths of the plasma, in this case the Debye length: $k\lambda_D \ll 1 \Leftrightarrow \omega/|\mathbf{k}| \gg v_t$. If $\omega \geq \omega_p$ the wave can propagate through the plasma. If $\omega < \omega_p$ the wave is partially reflected and partially damped inside the plasma. In this second case, the wave grazes the plasma surface for a length in the range of a *skin depth*, given by $d = c/\sqrt{\omega_p^2 - \omega^2}$, and is exponentially damped on larger length scales. Note that for $\omega \ll \omega_p$ the skin depth becomes a property of the plasma itself $d \approx c/\omega_p^2$. It is useful to introduce the *critical density*, namely the density at which the plasma frequency would be equal to the wave frequency:

$$n_c = \frac{m_e \omega^2}{4\pi e^2} = 1.1 \times 10^{21} \text{ cm}^{-3} \left(\frac{\lambda}{1 \mu\text{m}} \right)^{-2},$$

where the plasma frequency has been approximated with the electron plasma frequency and $\lambda = 2\pi c/\omega$ is the wavelength. The above considerations on the frequencies can be transferred into considerations on the densities. Given a plasma and a monochromatic wave, three different regimes can be identified:

- $n_e < n_c(\omega)$ *under-critical plasma*: the wave can propagate through the plasma,
- $n_e > n_c(\omega)$ *over-critical plasma*: the wave cannot propagate through the plasma,
- $n_e \approx n_c(\omega)$ *near-critical plasma*: intermediate situation where complex mechanisms arise and a strong coupling between wave and plasma is observed (see [12], [13]).

Misusing these simple considerations is useful for a preliminary classification of the physical processes that characterize the interaction between a laser pulse and a plasma (the above is only true for a plane wave in linear theory, even though provides a good description also for laser-plasma interaction). As an example, for a laser pulse with $\lambda = 0.8 \mu\text{m}$ wavelength the critical density is $n_c = 1.7 \times 10^{22} \text{ cm}^{-3}$. Given that the electron density of a fully ionized solid is typically about $n_e \sim 10^{22} \text{ cm}^{-3}$, it follows that the interaction between such lasers and a solid would occur in a strongly over-critical regime ($n_e \sim 200n_c$). These concepts are not trivially extended to relativistic laser-plasma interaction because of its inherent nonlinearities. Qualitatively, relativistic effects can be taken into account substituting m_e with $\gamma_e m_e$, where the (averaged of a laser cycle) relativistic factor is given by $\gamma_e = \sqrt{1 + a_0^2/2}$ (see Gibbon [14], Macchi [15]). Hence, the critical density increases for a given ω becoming:

$$n_c^{rel} = \gamma_e \frac{m_e \omega^2}{4\pi e^2}.$$

If $\gamma \gg 1$, a wave that cannot propagate in the non-relativistic limit may instead propagate in the relativistic case if $n_c < n_e < \gamma n_c$. This is effect is called *relativistic self-induced transparency*. If a pulse with $\lambda = 0.8 \mu\text{m}$ is ultra-intense, say $I = 10^{20} \text{ W/cm}^2$, then $a_0 = 6.8$ and $n_c^{rel} = 8.3 \times 10^{22} \text{ cm}^{-3}$. Depending on the specific solid density, relativistic transparency effects may arise.

1.4.1 Laser-particle interaction

The complex, nonlinear physics of laser-matter interaction can be coarsely simplified by considering one single particle. Let the electric field representative of the laser pulse be $\mathbf{E}(\mathbf{x}, t) = \mathbf{E}_0(\mathbf{x}) \cos(\mathbf{k} \cdot \mathbf{x} - \omega t)$, where \mathbf{k} is the propagation direction, ω is the pulse frequency and $\mathbf{E}_0(\mathbf{x})$ is the non-constant field amplitude. A non-relativistic particle under the action of such field moves according to $\ddot{\mathbf{x}} = \frac{q}{m} (\mathbf{E} + \frac{1}{c} \dot{\mathbf{x}} \times \mathbf{B})$, which can be solve in an approximated fashion. Using a perturbative approach, the particle trajectory is decomposed into two contributions as $\mathbf{x} = \mathbf{x}^{(0)} + \mathbf{x}^{(1)}$. The zero order dynamic $\mathbf{x}^{(0)}$ is the solution of the linearised version of the equation of motion. Expanding the electric field at

first order around the initial position of the particle \mathbf{x}_c as $\mathbf{E} \simeq \mathbf{E}(\mathbf{x}_c) + [(\mathbf{x} - \mathbf{x}_c) \cdot \nabla] \mathbf{E}|_{\mathbf{x}=\mathbf{x}_c}$, linearising yields to:

$$\mathbf{x}^{(0)} = \mathbf{x}_c - \frac{q\mathbf{E}_0(\mathbf{x}_c)}{m\omega^2} \cos(\mathbf{k} \cdot \mathbf{x}_c - \omega t) = \mathbf{x}_c - \frac{\mathbf{v}^{(0)}}{\omega} \cos \phi$$

After some computations, the following perturbative result can be derived:

$$\ddot{\mathbf{x}}^{(1)} = -\left(\frac{q}{m\omega}\right)^2 [(\mathbf{E} \cdot \nabla)\mathbf{E} + \mathbf{E}_0 \times \text{curl } \mathbf{E}_0 \sin^2 \phi + (\dots) \sin \phi \cos \phi]. \quad (1.18)$$

A force acting on the charge that only depends on the non-uniformities in the spatial profile of the field exists. Two terms can be identified: an *oscillating* force and a *secular* force. The former vanishes when averaging the equation over one laser oscillation period $T = 2\pi/\omega$ as opposed to the latter. Averaging the equation on the oscillation period (denoted with $\langle \cdot \rangle$) to only take into account the secular component yields to:

$$m \frac{d^2 \langle \mathbf{x}^{(1)} \rangle}{dt^2} = -\frac{q^2}{4m\omega^2} \nabla |E_0|^2 = \mathbf{F}_p \quad (1.19)$$

where \mathbf{F}_p is the so-called *ponderomotive force*. Considering nonlinearities in this approximated fashion leads to an effective force that acts because of the non-uniformities in the spatial profile of the field. The ponderomotive force pushes the particles towards regions with smaller ponderomotive potential $U_p = \frac{q^2}{4m\omega^2} |E_0|^2$ (see figure 1.4).

For high intensities the relativistic effects cannot be neglected. In this case, the ponderomotive force and potential have to be corrected according to the following (see Mulser and Bauer [16] or [17]):

$$\mathbf{F}_p = -mc^2 \nabla \sqrt{1 + \langle \mathbf{a}^2 \rangle}, \quad (1.20)$$

where $\mathbf{a} = e\mathbf{A}/mc^2$ and \mathbf{A} is the vector potential defined by $\mathbf{E} = -(1/c)\partial\mathbf{A}/\partial t$. Note that \mathbf{F}_p is inversely proportional to the mass, hence its effect is much stronger on electrons than on protons and ions. Although this argument is only valid for a single particle in interaction with a laser pulse, it can suggest what happens when the interaction is with matter. Qualitatively, if the action of \mathbf{F}_p occurs on times scales smaller than the characteristic time scales of the ion dynamic, then electrons are forced to separate from ions and matter ionizes. Therefore the ponderomotive force is a source of charge separation within matter; in turn charge separation is a source for high magnitude, longitudinal electric fields that ultimately can be exploited to accelerate particles.

The ponderomotive force arising from the interaction between a laser pulse and a charged particle possibly represents the simplest concept related to the physics of laser-plasma interaction. Despite all the simplifications, it allows to build qualitative arguments to understand why it is possible to accelerate particles using plasmas. Consider an ultra-intense, ultra-short laser pulse irradiated onto a plasma.




Figure 1.4: Schematic action of the ponderomotive force.

Suppose that such pulse is able to propagate through the considered plasma for long distances (even \sim cm), i.e. the interaction happens in the under-critical regime. Travelling across matter, the pulse acts on the electrons via the ponderomotive force. The electrons are pushed away from the region where the pulse is located towards regions where the (squared) amplitude of the field is lower (see figure 1.4). Therefore, electrons that have not met the pulse yet are moved forward along the pulse propagation direction. In the meantime, the pulse moves further at the speed of light and is able to overtake those electrons. At this point the same electrons are pushed backwards, again via the ponderomotive force. Essentially, the electrons are pushed back and forth by the laser. This kind of interaction can trigger a collective oscillation of the electrons, i.e. a plasma wave, consisting in a wake of oscillating electrons that is produced behind the propagating pulse. Furthermore, for proper laser wavelengths, the oscillation may be resonant. This phenomena is called *wake-field generation* and is crucial for laser-based electron acceleration (for more insights on this topic refer to [18]). Indeed, if resonance occurs, huge electric fields arise within the plasma increasing up to a maximum value (about GV/cm) at which the so-called *wavebreaking* occurs. In this situation part of the oscillating electrons are trapped in the regions close to the maximum of the field and are effectively accelerated. See figure 1.5 to visualize this effect.




Figure 1.5: Schematic wakefield acceleration. Taken from [19].

Suppose now that the laser pulse is irradiated onto a relatively thin plasma that behaves like a sort of mirror with respect to the laser pulse, i.e. the interaction happens in the over-critical regime. Typically, a solid material is needed for this to happen. In this case the

interaction is superficial and the pulse is reflected backwards. During the interaction the laser pulse is able to give a fraction of its energy to the electrons populating the front side of the plasma through many different mechanisms. Among them, if the laser is P polarized, the oscillating part of the force in (1.18) tears off the electrons from the solid into vacuum. Such electrons are re-injected in the plasma at very high energies, pass through the thin target and ultimately generate a strong charge separation, i.e. strong electric fields on the rear side. This field can ionize the hydrocarbon contaminants which are always present on the back face of the target (unless special care is taken), which can then be accelerated by the electric field up to high energies. A more detailed description of laser-induced ion acceleration follows in the next section.

1.5 Laser-driven ion acceleration

In the most typical laser-driven ion acceleration experiments a laser pulse is irradiated onto a target, for example a thin solid foil, that rapidly turns into plasma. The direct interaction between the pulse and the electron population generates a collective displacement of charge that in turn produces strong electric fields. Such fields accelerate the ions, that try to restore neutrality following the electrons in a ballistic way (see figure 1.6). In the year




Figure 1.6: Artistic depiction of a typical laser-driven ion acceleration experiment. From [20].

2000 three experiments of superintense laser-plasma interaction demonstrated the possible generation of multi-MeV protons emitted from the non-irradiated side of μm -thick solid targets [21, 22, 23]. Since then great interest has been addressed towards the investigation of laser-based ion sources both for the applicative potential and the research of new regimes of radiation-matter interaction (see [20, 24] for a review). Many diverse experimental configurations have been tested in different laser facilities worldwide, varying both target characteristics (material, thickness, density, structure, morphology,...) and laser properties

(intensity, power, duration, polarization, angle of incidence,...). Some of the recorded maximum proton energies are shown in figure 1.7 as functions of the pulse power. Low repetition rate (few shots per day) laser pulses with high energy and high power (160-200 J, 500-800 fs) have been recently used to produce ion bunches with cutoff energies up to 85 MeV and 10^9 particles per bunch with state-of-the-art laser facilities, which is the current record as regards the maximum proton energy [3].




Figure 1.7: Experimental scaling of the maximum proton energy as function of the pulse power. Data refer to different experimental campaigns performed in the facilities according to the legend. Left: ultra-short pulses with durations around 30 fs. Right: "long" pulses with durations of hundreds of fs; the big red dot represents the current record.

One of the most interesting features of multi-MeV ions is that, when travelling through dense matter, they deliver most of their energy at the end of their path. The energy absorbed per unit mass by the crossed medium as a function of the travelled path length is characterized by a pronounced peak right before the particles come to rest, known as *Bragg peak* (see figure 1.8). Physically, the charged particles ionize the atoms of the material as they move across it, gradually losing energy. The peak occurs because the interaction (Coulomb collisions) cross section strongly grows as the particle energy decreases, hence the halting process becomes more and more efficient. The same does not happen for electrons, x rays and γ rays. This kind of energy deposition is suitable for highly localized energy deposition, desirable in many applications, e.g. hadron therapy [20, 25]. The perspective of compact laser-based ion sources for hadrontherapy has been one of the main drivers of high intensity laser-matter interaction research. The aforementioned ion acceleration process produces naturally broad exponential energy spectra, which are unsuitable for direct applications in hadrontherapy. Moreover, the maximum energy of laser-accelerated ions is still too low to be of medical interest. These issues have motivated the research on other ion acceleration mechanisms and/or advanced target designs. Another promising option, which has attracted some attention recently, is the possibility to use laser-driven ion sources for other applications with relaxed requirements, such as production of radionuclides [26]

or production of neutrons [27].




Figure 1.8: Example of the depth dependence of the relative dose, i.e. the energy absorbed per unit mass, for different radiations in water. Proton and C ion profiles are characterized by the Bragg peak at the end of the path.

In 3 some experimental data, especially regarding the accelerated ions, will be compared with numerical simulation results. For this reason it is useful to mention that experimental measurements of the properties of the accelerated ions require specific diagnostics. Among them, the use of *radiochromic films* (RCF) is well-established since the early experiments of laser-driven ion acceleration by Snavely [23]. Simplifying, a RCF is a layer of radiation-sensitive material that turns from nearly transparent to a shade of blue upon interaction with ionizing radiation. The darkness of the shade is proportional to the absorbed dose, hence it gives information on the flux of ions travelling towards the layer. In ion acceleration experiments usually RCFs are arranged in stacks, so that each RCF acts as an energy filter for the successive ones. The accelerated protons travel through the stack and come to rest at a certain depth, corresponding to their Bragg peak. Hence, the signal on a given layer is only due to those ions with energy equal or greater than the Bragg peaks energies falling within the very same layer. An example of RCF stack is shown in figure 1.9. Very complex structures are observed at different energies and with different materials. Simulated RCFs can be produced from simulations output; an analysis of numerical RCFs with foam-attached targets will be presented in chapter 3. RCFs represent only one of the ion diagnostics used in experiments. Other kind of film detectors can be employed, such as CR39 layers, with the advantage of being insensitive to x rays and electrons. Moreover, deeper information on the accelerated beam spectrum can be obtained using for example Thomson parabola spectrometers, based on the same idea of mass spectrometry.




Figure 1.9: Example of RCF stack obtained with the TRIDENT short-pulse laser at the Los Alamos National Laboratory. Indicated beneath each film is the proton Bragg peak energy. Reprinted from [28].

1.5.1 Acceleration mechanisms: target normal sheath acceleration

Several mechanisms have been introduced to interpret the accelerated ions properties as functions of the laser and target characteristics. Among them, the Target Normal Sheath Acceleration (TNSA) mechanism first proposed by Wilks *et al.* [29] has proven to be the most solid, since most of the experiments can be (at least partially) explained according to this scheme.

The basic concept is the following. Electrons have much lower inertia if compared to ions, since the elementary mass is about three orders of magnitude smaller than the proton mass ($m_p \approx 1836m_e$). This implies that, at today feasible intensities (up to 10^{21} W/cm²), the impinging laser pulse directly interacts only with the electron population (ions would be directly accelerated by the laser pulse only at intensities about 10^{24} W/cm², but this is far beyond present-day capabilities). Thereupon, many electrons on the front side of the target – called *fast* or *hot electrons* – absorb laser energy through different non-collisional mechanisms (e.g. J×B heating and vacuum heating, see Gibbon [14]) and gain enough kinetic energy to move through the bulk and escape from the rear side, where they form a cloud spreading for few λ_{De} . Hence a space-charge displacement occurs which gives rise to a strong longitudinal (along target normal direction) electric field, called *sheath field*, rapidly decaying outside the target after few μ ms. This process leads to ion acceleration along the target normal direction. Protons coming from the ionization of the impurities located on the rear surface are the firsts to be accelerated up to tens of MeVs, because of both their favourable position and small charge-to-mass ratio. Other kinds of ions may be accelerated up to tens of MeVs per nucleon on longer time scales.

Rough estimations of the sheath field and of the maximum ion energy can be obtained using simple dimensional arguments assuming the only relevant parameters to be the size of the hot electron cloud L_c , the hot electron temperature T_h and density n_h . The sheath field can be estimated as:

$$E_s \sim \frac{T_h}{eL_c}.$$

Now, T_h can be estimated using the ponderomotive scaling assuming that hot electrons kinetic energy is of the order of that in the electric field of the laser in vacuum, averaged over one oscillation:

$$T_h \sim \mu_h m_e c^2 (\gamma - 1) = \mu_h m_e c^2 \left(\sqrt{1 + \frac{a_0^2}{2}} - 1 \right),$$

where μ_h is the fraction of laser energy absorbed by the hot electron population. The sheath size L_c is of the order of the electron Debye length, thus:

$$L_c \sim \lambda_{Dh} = \sqrt{\frac{T_h}{4\pi e^2 n_h}}.$$

Considering typical experimental values of the parameters such as a laser irradiance of $I_0 \lambda^2 \sim 10^{20}$ W/cm² μm^2 (i.e. $a_0 \sim 8.5$), a solid-density plasma $n_h \sim 10^{23}$ cm⁻³ and $\mu_h = 0.1$, one obtains a hot electron temperature of \sim MeV, a sheath size of $\sim \mu\text{m}$ and a sheath field of ~ 1 MV/ μm . The energy acquired by a test ion in the sheath field is

$$E \sim ZeE_s L_c \sim ZT_h$$

which results in MeVs and in a one-half power scaling with the intensity: $E \sim \sqrt{I}$. More sophisticated scaling laws can be derived using different theoretical models for TNSA mathematical description (see [30]).

As a general consideration, the sheath field can be increased either by increasing the hot electron temperature or density. Therefore, the TNSA mechanisms can be optimized by properly tuning these quantities. One option is to enhance the laser energy absorption by suitably structuring the target, as explained in the next section.

Lastly, as already mentioned, many different acceleration mechanisms arise during the interaction of high intensity laser pulses and matter, such as radiation pressure acceleration (RPA), collisionless shock acceleration, Coulomb explosion, hole boring [15]. Each of these may dominate the physics of the interaction depending on the system configuration. For example the effects of RPA should be strong when using very intense circularly polarized laser pulses irradiating very thin targets (tens of nm). Anyhow, a TNSA component of the process is commonly observed, even when trying to suppress its effects with different strategies when studying alternative acceleration mechanisms.

1.5.2 Nanostructured targets

The actual employment of laser-based ion sources requires further optimization of many parameters, implying an enhancement of the acceleration scheme. In particular, higher numbers of accelerated particles with higher cutoff energies are needed. Moreover it will




Figure 1.10: Schematic target normal sheath acceleration. The laser-target interaction produces hot electrons at the front side that move across the bulk and reach the rear side where they form a sheath in vacuum. The electric sheath field generated by the charge separation is almost perpendicular to the rear surface and is able to accelerate ions.

be necessary to find new relaxed experimental configurations to be able to produce high-energy ion beams at high repetition rate (at least 10 Hz) with relatively moderate laser intensities in a controlled process. Improvements can be achieved through the design of non conventional targets.

Recent experiments have shown an increase in the total number of accelerated ions as well as in their maximum energy using special double-layer targets with respect to bare solid foils (see [31], [32], [33], [34]). A schematic representation of both kinds of targets is shown in figure 1.11. The idea is to induce an enhanced TNSA-like mechanism, increasing laser absorption and fast electron generation by attaching an additional layer on the illuminated side of the solid target (most certainly the addition of such layer results in an extra electron population, absent in bare solid targets). Besides improving the acceleration process, the actual interest in this kind of targetry also comes from the non-trivial advantage of their *robustness*, since they can be easily handled without risking their damage as opposed to other kinds of advanced targets used to improve the process as well, such as ultra-thin nm-thick foils for enhanced RPA (see [35, 36]). Therefore they represent a promising solution towards the repetitive regime, necessary for the generation of ion *beams*. The density of such layer should be around the plasma critical density in order to take advantage of the efficient coupling between laser and matter expected at near-critical density. In this scenario, both surface and volume laser absorption mechanisms arise, while with simple solid targets only the former takes place. Hence the presence of the additional layer leads to an enhanced generation of relativistically hot electrons. Since typical values of the laser pulse wavelength used in high-intensity experiments are $\lambda \sim 0.8\text{-}1 \mu\text{m}$, the corresponding

critical density is

$$n_c = \frac{m_e \omega^2}{4\pi e^2} = 1.1 \times 10^{21} \text{ cm}^{-3} \left(\frac{\lambda}{1 \mu\text{m}} \right)^{-2} \sim 0.88 - 1.1 \times 10^{21} \text{ cm}^{-3},$$

which implies mass densities of few mg/cm^3 , few times the density of air. The manufacturing of such materials is very challenging because these density values are challenging for gas jets and too low for solids. One (almost the only) way to produce the additional layer – called *foam* hereinafter – attached to a solid foil is to grow a porous nanostructured carbon layer directly on a solid substrate by means of the *pulsed laser deposition* technique [37]. In a gas-filled deposition chamber a laser pulse is irradiated on a carbon target that is consequently ablated, producing a cloud of vaporized species. The ablated species expand and assemble in aggregates that are then deposited on the solid substrate. This procedure leads to the production of a "thin" low-density carbon film, coating the substrate. The resulting material is made by small (tens of nanometres), dense (about solid carbon density) clusters aggregated in structures on the micrometer scale (see figure 1.13). Such growing process gives rise to very open materials with low mean densities on the mesoscale (see figure 1.14). The filling factor, i.e. the volume fraction actually occupied by the solid density carbon, is about 0.003-0.005. By properly tuning the set of deposition parameters it is possible to obtain foams with tailored density, thickness, morphology and structure as shown in figure 1.12. From the cross sections in figure 1.12 a cauliflower-like growth is evident at the lower pressures, while more disordered configurations are revealed at higher pressures. In any case, as a consequence of the manufacturing process, all these materials are inherently nanostructured and exhibit complex spatial non-uniformities. A good even if simplified mathematical model for such processes is given by the diffusion-limited aggregation model, as will be explained later in section 3.1.




Figure 1.11: Schematic sketch of the targets. Conventional targets are overdense, μm -thick solid foils (left panel). Foam-attached targets are special double-layer targets where a near-critical foam layer is placed in front of a μm -thick solid foil to enhance the acceleration process (right panel).

The behaviour of this kind of targets has been experimentally tested both at moderate laser intensities and in the fully relativistic regime [32, 33]. A significant improvement in




Figure 1.12: Scanning Electron Microscopy images that show how the morphology of the foam targets depend on the ambient gas properties. Top views and cross sections alternate on the rows.

the performance has been observed with respect to the bare solid target case, with gain factors (maximum proton energy with foam/maximum proton energy with simple foil) for the maximum proton energy in the ranges 2-3 and 1.4-3 in the two intensity regimes respectively. Figure 1.15 shows the experimental proton energy spectra with and without the foam in different situations at relativistic intensities ($I \sim 10^{20} \text{ W/cm}^2$).

The ongoing H2020 ERC project ENSURE ("Exploring the New Science and engineering unveiled by Ultraintense ultrashort Radiation interaction with mattEr") at the Micro and Nanostructured Materials Lab, Department of Energy, Politecnico di Milano, focuses on the experimental and numerical investigations of laser-driven ion acceleration with the nanostructured, foam-attached targets described above. The present thesis work has been carried out within the framework of the same project and concerns the simulation of the physical systems described along this chapter, i.e. nanostructured plasmas irradiated by high intensity laser pulses.




Figure 1.13: High Resolution Scanning Transmission Electron Microscopy image of a foam sample. The building blocks of the aggregate are nanoparticles of $\sim 10\text{-}20$ nm size. Reprinted from [37].




Figure 1.14: Scanning Electron Microscopy images showing the aggregates assembled on the micrometer scale. Reprinted from [37].




Figure 1.15: Energy spectra of the protons obtained with a laser intensity of $4.1, 3.5, 3.7 \times 10^{20}$ W/cm² for S, P and C polarization respectively. The targets are double-layer targets with $8 \mu\text{m}$ -thick foam (spectra a, b, c respectively) and single targets (spectra d, e and f). The spectra are collected along the target normal direction. The inset shows the electron energy spectra when using double-layer targets with $12 \mu\text{m}$ -thick foam with S, P and C polarization, (spectra a, b, c respectively) and using single targets for S and C polarization (spectra d and f). Reprinted from [33].

Chapter 2

Numerical methods for the Vlasov-Maxwell system

At present day, mathematical models represent an essential tool for the description of physical systems in many scientific fields. Often, these models are way too complex to be solved analytically, thus numerical schemes for their approximation must be considered. This also holds within the framework of collisionless kinetic plasma physics, where the governing equations are given by the Vlasov-Maxwell system (1.17). These are strongly nonlinear, non-local, coupled partial differential equations to which a numerical solution is desirable. For their approximation, two main classes of numerical schemes exist: *phase space-grid based methods* and *particle methods*. A brief overview of the former is presented in section 2.1. Section 2.2 contains a detailed description of the particle-in-cell method, the most exploited among particle methods and the one strictly related to the present thesis work. Advantages and disadvantages of both methods are summarized in section 2.3. At the end of the chapter, section 2.4 details the motivations and aims of this thesis work.

2.1 Phase space-grid based methods

In this section we present the ideas underlying some of the most exploited methods for Vlasov equation approximation that use a grid of phase space. The distribution function is discretized on a grid and at each time step is evolved according to Vlasov equation (1.16), while the electromagnetic field is computed separately with some standard algorithm once f is known. Codes that implement this kind of methods are often called *Vlasov codes*.

Since the distribution function is defined on a six-dimensional space, it is easy to ascertain that these methods require a lot of computational resources in more than one dimension.

Consider a 3D3V simulation, i.e. in a full six-dimensional phase space, on a grid made of 1000 points in every direction. Overall, there are 10^{18} nodes. Hence, the storage of the distribution function in double precision would require 8 ExaBytes of RAM, far beyond present-day supercomputer capabilities.

For this reason grid-based methods are mainly used for low-dimensional, reduced models simulations. The simplest example is given by the 1D and 2D electrostatic Vlasov equation for one species in a neutralizing background (here all physical constants are set to 1 for the sake of simplicity):

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{x}} f + \mathbf{E}(\mathbf{x}, t) \cdot \nabla_{\mathbf{v}} f = 0,$$

coupled with Poisson equation:

$$-\Delta\phi = 1 - \int f d\mathbf{v},$$

where ϕ is the electrostatic potential defined by $\mathbf{E} = -\nabla\phi$. Another reduced model is the 4D drift-kinetic model, used especially for highly magnetized plasma in tokamaks. Assuming a large and uniform magnetic field \mathbf{B} the following equation can be derived

$$\frac{\partial f}{\partial t} + \left(\mathbf{v}_{\parallel} + \frac{\mathbf{E}(\mathbf{x}, t) \times \mathbf{B}}{B^2} \right) \cdot \nabla_{\mathbf{x}} f + \mathbf{E}_{\parallel} \cdot \nabla_{\mathbf{v}} f = 0,$$

where \parallel denotes the magnetic field direction. Again, the equation is coupled with Poisson equation. One last example, useful for nonlinear interaction of high intensity ultra short laser pulses with plasmas, consists in the relativistic Vlasov equation (1.16) where f , \mathbf{E} and \mathbf{B} depend only on one or two space variables and f depends on one, two or three momentum components. A further reduction can be applied depending on the wave polarization and/or assuming a specific form for the solution (see [38]). In all these cases Vlasov equation can be written as an advection equation where the transport field is divergence free.

A common problem suffered by grid-based methods is that they are affected by the inability to fully resolve phase space filamentation phenomena (see 1.3.2). Indeed, filamentation often occurs during the evolution of the distribution function and, at some point during the computation, cannot be further resolved. Sooner or later, the grid will become coarse enough not to resolve such small scales. This mechanism results in numerical instabilities and non-physical oscillations that may lead to nonpositive values of the distribution function, making long time simulations unreliable. To dispose of the filamentation it is possible to add artificial dissipation or, alternatively, Fourier filter large k-modes or smooth gradients in the phase space. Nevertheless, grid-based methods allow a uniform and fine phase space resolution irrespective of the magnitude of f .

2.1.1 Time splitting

Many grid-based schemes split Vlasov equation into two advection equations, as originally proposed by Cheng and Knorr [39] for the electrostatic case and later by Cheng for the magnetized case [40]. The distribution function is evolved from time t to time $t + \Delta t$ in four steps as follows. First solve for half a time step

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{r}} f = 0,$$

with fixed \mathbf{v} . Second, solve the field equations (1.7). Third, integrate for a whole time step

$$\frac{\partial f}{\partial t} + \frac{\mathbf{F}(\mathbf{r}, t)}{m} \cdot \nabla_{\mathbf{v}} f = 0,$$

with fixed \mathbf{r} . Lastly, solve again the first equation for half a time step. In the first equation the advection field is independent from the advection variable, so it can be solved explicitly. The same also holds for the second equation only under additional hypothesis, for example in the electrostatic case.

2.1.2 Semi-Lagrangian methods

Semi-Lagrangian methods are based on the conservation property of the distribution function along the characteristics. For the sake of simplicity, a compact notation is used: $\mathbf{x} = (\mathbf{r}, \mathbf{v})$ summarizes the position and velocity Eulerian coordinates, while $\mathbf{a}(\mathbf{x}, t)$ is the total advection field. Thus Vlasov equation can be written as an abstract advection equation as

$$\frac{\partial f}{\partial t} + \mathbf{a}(\mathbf{x}, t) \cdot \nabla_{\mathbf{x}} f = 0$$

and the characteristics equation (1.13) becomes

$$\frac{d\mathbf{X}}{dt} = \mathbf{a}(\mathbf{X}(t), t).$$

The unique solution of this problem with initial condition $\mathbf{X}(s) = \mathbf{x}$ is indicated with $\mathbf{X}(t; \mathbf{x}, s)$.

Backward method The classical method consists in the backtracking of the characteristics: their path is followed backwards at each time step, then the distribution function is reconstructed on a grid. Consider a grid on phase space $\{\mathbf{x}_i\}$ and a partition of the time interval in steps $\{t^k\}$. Assume the problem is solved at time step t^n . The algorithm used to update the distribution function from time step n to $n+1$ is made of two steps. First, find the origins of the characteristics ending at each grid point, i.e. compute $\mathbf{X}(t^n; \mathbf{x}_i, t^{n+1})$. Second, use the conservation along the particle trajectories

$$f(\mathbf{x}_i, t^{n+1}) = f(\mathbf{X}(t^{n+1}; \mathbf{x}_i, t^{n+1}), t^{n+1}) = f(\mathbf{X}(t^n; \mathbf{x}_i, t^{n+1}), t^n),$$

to approximate the distribution function on the grid node \mathbf{x}_i at time t^{n+1} . Since in general $\mathbf{X}(t^n; \mathbf{x}_i, t^{n+1})$ does not coincide with any grid point, $f(\mathbf{x}_i, t^{n+1})$ is computed by interpolation from the values on the grid.

In general a numerical algorithm is required to perform the first step, e.g. fixed point or Newton method. Yet in many situations time splitting between position and velocity advections can represent an efficient way to solve the problem, with the advantage of allowing the explicit computation of the feet of the characteristics. However, splitting leads to not conservative equations when the advection fields of the split equations are not divergence free. This in turn may lead to large errors and un-physical behaviours, especially when phase space filamentation occurs in nonlinear simulations. In this case the total number of particles is strongly not conserved,

For the second step, to avoid too much dissipation high order interpolation is needed. Typically, cubic splines [39] or cubic Hermite with derivative transport interpolation [41] schemes are used.

Forward method The general idea is again to use the conservation of the distribution function along particle trajectories, but here the characteristics are followed forward in time, not backwards. The grid points are advanced in time as if they were plasma particles and for each time step the distribution function is reconstructed on a grid using a suitable convolution kernel. The use of cubic B-splines to scatter the values of f on the grid offers a good compromise between efficiency and accuracy. This method can be thought as a particle method where f values are deposited on the whole phase space, not only on the configuration space. A hybrid method where the reconstruction is performed periodically was proposed in [42].

Conservative method In this method the unknown is the average of the distribution function over one cell $\frac{1}{V} \int_V f(\mathbf{x}) d\mathbf{x}$, which is another quantity conserved along the characteristics. Starting from the conservative form of the Vlasov equation, time splitting applied to separate all the direction yields to 6 one-dimensional advection equations that can be written in conservative form as $\partial f / \partial t + \partial(a(x, t)f) / \partial x = 0$. This method consists of three steps:

1. reconstruction: construct on every cell a high order polynomial function whose cell average is equal to the cell average of f at time step n , f_i^n , for example using an interpolation method;
2. resolution: compute the origins of the characteristics ending at the grid points, i.e. backtrack each cell, solving the equation of the characteristics;

3. projection: compute cell average f_i^{n+1} at the new time step $n + 1$ using the conservation property:

$$\int_{x_{i-1/2}}^{x_{i+1/2}} f^{n+1}(x) dx = \int_{X(t^n; x_{i-1/2}, t^{n+1})}^{X(t^n; x_{i+1/2}, t^{n+1})} f^n(x) dx,$$

where $f^n(x)$ is the high order reconstruction found in step 1.




Figure 2.1: Sketch of the idea underlying semi-Lagrangian methods. Backwards (left), forward (middle) and finite volume (right). Taken from [43].

2.1.3 Spectral methods

Different kinds of transform methods have been proposed for special applications. A Fourier-Fourier transformed version of the splitting algorithm was introduced in [44, 45] for the electrostatic 1D1V case. A Fourier transformation is applied both in space and velocity to the Vlasov-Poisson system. Each step of the modified splitting algorithm consists in performing a forward fast Fourier transform, the application of a phase shift to the expansion coefficients and a backward Fast Fourier transform. In addition, a special filtration of high frequency modes in the velocity transform is carried out to avoid phase space filamentation.

A 2D2V and 3D3V spectral method are presented in [46] and [47] respectively. The Vlasov equation is Fourier transformed in the velocity space, not in configuration space. The obtained equation is numerically solved. Suitable outflow boundary conditions in the transformed velocity space are used to dispose of the filamentation phenomena.

Other types of spectral methods using different transformation procedures have been studied in [48, 49, 50, 51, 52].

However, transform methods often do not represent the most advantageous choice because of some inherent difficulties, such as code parallelization, boundary conditions definition and prevention of the filamentation.

2.2 The particle-in-cell method

The *particle-in-cell* (PIC) method is the most widely used and well-established numerical scheme for the approximated solution of the Vlasov-Maxwell system. The idea of such method was conceived between the late fifties and early sixties notably by John Dawson and Oscar Buneman in their pioneering works on the first computer experiments in plasma physics [53, 54, 55, 56]. A large literature on this method exists, concerning both the physical and mathematical point of view. For the former aspect refer to Birdsall and Langdon [57] or Hockney and Eastwood [58], for the latter to [59, 60, 61, 62].

2.2.1 Basic principles of the PIC scheme

The basic idea of the PIC scheme is to adopt a combined Lagrangian-Eulerian approach representing the distribution function f by a set of Lagrangian particles and computing the electromagnetic field on a finite-dimensional Eulerian grid. To be more specific, the distribution function of each plasma population is sampled using a collection of computational particles – called *macro-particles* – each representative of a high number of real particles. The current density field associated to the motion of such particles is scattered on a mesh on the configuration space, on which the electromagnetic field is then computed. From the grid values of the electromagnetic field, the Lorentz force acting on every particle is calculated by interpolation, hence particles motion can be updated. The procedure is iterated for as many times as desired (see figure 2.2).

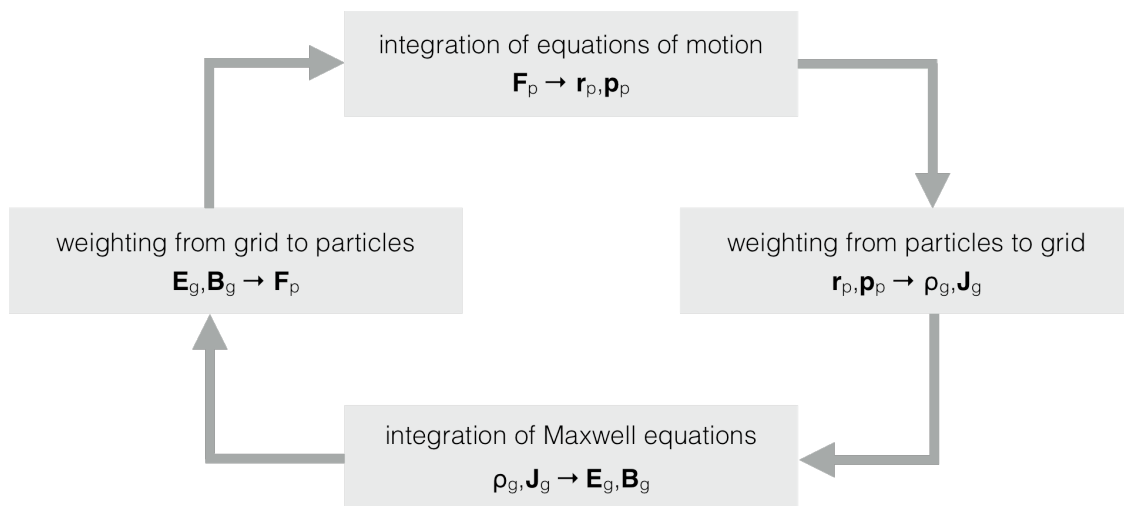


Figure 2.2: One time step in a particle-in-cell simulation program. Adapted from [57].

In the continuous description the distribution function $f(\mathbf{r}, \mathbf{p}, t)$ represents the density of particles at time t in a neighbourhood of the point (\mathbf{r}, \mathbf{p}) of the phase space. In the discrete

PIC description the concept of distribution function is replaced with the concept of macro-particles. Formally, the PIC idea corresponds to approximate the distribution function of each plasma species by a sum over the macro-particles of Dirac masses centred on each particle trajectory in phase space $(\mathbf{r}_p(t), \mathbf{p}_p(t))_{1 \leq p \leq N_p}$:

$$f(\mathbf{r}, \mathbf{p}, t) \approx f_h(\mathbf{r}, \mathbf{p}, t) = A \sum_{p=1}^{N_p} \delta(\mathbf{r} - \mathbf{r}_p(t)) \delta(\mathbf{p} - \mathbf{p}_p(t)), \quad (2.1)$$

where A is a normalization constant and N_p is the total number of particles of the considered species. The distribution function of every plasma population is approximated in this fashion.

In practice positions and momenta of the macro-particles are initialized as $\mathbf{r}_{p,0}$ and $\mathbf{p}_{p,0}$ so that $f_{h,0}(\mathbf{r}, \mathbf{p}) = A \sum_{p=1}^{N_p} \delta(\mathbf{r} - \mathbf{r}_{p,0}) \delta(\mathbf{p} - \mathbf{p}_{p,0})$ is an approximation of the initial distribution function $f_0(\mathbf{r}, \mathbf{p})$ in some sense. Given $f_0(\mathbf{r}, \mathbf{p})$, the initialization step is often performed using a Monte-Carlo strategy, namely the initial positions and momenta are randomly chosen according to the probability density associated to $f_0(\mathbf{r}, \mathbf{p})$. Once the initialization is completed, the particles are advanced according to the equations of motion:

$$\left\{ \begin{array}{l} \frac{d\mathbf{r}_p(t)}{dt} = \mathbf{v}_p(t) \\ \frac{d\mathbf{p}_p(t)}{dt} = q \left[\mathbf{E}(\mathbf{r}_p(t), t) + \frac{\mathbf{v}_p(t)}{c} \times \mathbf{B}(\mathbf{r}_p(t), t) \right] \\ \mathbf{r}_p(0) = \mathbf{r}_{p,0} \\ \mathbf{p}_p(0) = \mathbf{p}_{p,0}, \end{array} \right.$$

where $\mathbf{p}_p(t) = m\gamma_p(t)\mathbf{v}_p(t)$ and $\gamma_p(t) = \sqrt{1 + \mathbf{p}_p(t) \cdot \mathbf{p}_p(t)/m^2c^2}$.

It is easily verified that $f_h(\mathbf{r}, \mathbf{p}, t)$ is an exact solution of the Vlasov equation in the sense of distributions with initial condition $f_{h,0}$. Let $\varphi \in C_0^\infty(\mathbb{R}^3 \times \mathbb{R}^3 \times (0, \infty))$, then f_h defines the following distribution:

$$\langle f_h, \varphi \rangle = A \sum_{p=0}^{N_p} \int_0^\infty \varphi(\mathbf{r}_p(t), \mathbf{p}_p(t), t) dt.$$

Now, by definition and using the compactness of the support of the test function φ it

follows that

$$\begin{aligned}
\left\langle \frac{\partial f_h}{\partial t}, \varphi \right\rangle &= - \left\langle f_h, \frac{\partial \varphi}{\partial t} \right\rangle = -A \sum_{p=0}^{N_p} \int_0^\infty \frac{\partial \varphi}{\partial t}(\mathbf{r}_p(t), \mathbf{p}_p(t), t) dt \\
&= -A \sum_{p=0}^{N_p} \int_0^\infty \left(\frac{d}{dt} - \frac{d\mathbf{r}_p(t)}{dt} \cdot \nabla_{\mathbf{r}} - \frac{d\mathbf{p}_p(t)}{dt} \cdot \nabla_{\mathbf{p}} \right) \varphi(\mathbf{r}_p(t), \mathbf{p}_p(t), t) dt \\
&= - \left\langle \left(\frac{d\mathbf{r}}{dt} \cdot \nabla_{\mathbf{r}} + \frac{d\mathbf{p}}{dt} \cdot \nabla_{\mathbf{p}} \right) f_h, \varphi \right\rangle.
\end{aligned}$$

Actually (2.1) does not represent the ultimate approximation for the distribution function used in the PIC algorithm. Considering point-like particles in the configuration space leads to the onset of undesired numerical noise. To reduce numerical noise, the particles are allowed to have a finite extension in the configuration space, whereas are kept point-like in the velocity space (see figure 2.3). The approximated distribution function hence reads

$$f_h(\mathbf{r}, \mathbf{p}, t) = A \sum_{p=1}^{N_p} S(\mathbf{r} - \mathbf{r}_p(t)) \delta(\mathbf{p} - \mathbf{p}_p(t)), \quad (2.2)$$

where S – called *shape function* – is an even function localized in a neighbourhood of zero and normalized to unity ($\int S(\mathbf{r} - \mathbf{r}_p(t)) d\mathbf{r} = 1$) that describes the shape of the macro-particles, namely their charge distribution profile.




Figure 2.3: Sampling of the distribution function with macro-particles. Each macro-particle has a definite momentum, but is extended in space.

2.2.2 Theoretical derivation

The problem solved by the PIC scheme, i.e. the motion of the macro-particles, can be derived by substituting (2.2) in Vlasov equation (1.16). Since Vlasov and Klimontovich equations are identical from the mathematical point of view, the procedure is analogous to the proof that Klimontovich equation implies the equations of motion of the plasma particles (see appendix A for details). The only difference is that the Dirac delta functions in the momenta variable are replaced with the shape functions, which yields to a mean force defined by:

$$\tilde{\mathbf{F}}_p(t) = \int S(\mathbf{r} - \mathbf{r}_p(t)) \mathbf{F}(\mathbf{r}, \mathbf{v}_p(t), t) d\mathbf{r},$$

where $\mathbf{F}(\mathbf{r}, \mathbf{v}, t)$ is the Lorentz force. Summarizing, if the distribution function is replaced by its approximated expression in Vlasov equation, the relativistic equations of motion of the macro-particles can be obtained for all $p = 1, \dots, N_p$ and for all species:

$$\begin{cases} \frac{d\mathbf{r}_p(t)}{dt} = \frac{\mathbf{P}_p(t)}{\gamma_p(t)m} \\ \frac{d\mathbf{P}_p(t)}{dt} = \tilde{\mathbf{F}}_p(t) = q \int S(\mathbf{r} - \mathbf{r}_p(t)) \left[\mathbf{E}(\mathbf{r}, t) + \frac{\mathbf{P}_p(t)}{\gamma_p(t)mc} \times \mathbf{B}(\mathbf{r}, t) \right] d\mathbf{r}. \end{cases} \quad (2.3)$$

Every macro-particle moves as a particle of mass m and charge q under the action of the Lorentz force averaged over its volume. The time evolution of the distribution function corresponds to the particles motion in the phase space. This justifies the interpretation of f in terms of the macro-particles.

System (2.3) could be also deduced from Klimontovich equation replacing spatial Dirac delta-functions with shape functions in the definition of the microscopic distribution function. So, what does a PIC code really solve? Vlasov-Maxwell system, describing the average behaviour of a collisionless plasma, or the Klimontovich-Maxwell system, describing the exact behaviour of a real plasma? This is a non-trivial question, still unanswered in the community. Thus far it is still not clear if f_h is an approximation for the actual distribution function or for the microscopic distribution function. Indeed, both problems are strictly related to the PIC method, nevertheless many differences exist between PIC plasmas, real plasmas and Vlasov-Maxwell plasmas (see figure 2.4). In particular, two crucial approximations are made to numerically model a plasma in a PIC description. The first consists in the grouping of a large number (up to 10^{10}) of real particles into one single macro-particle, a fact known as *coarse-graining*; the second is the grid-based fields discretization. Furthermore, real plasmas are affected by collisions, while only long-range interactions characterize the behaviour of a Vlasov-Maxwell plasma. Considering all these issues makes it hard to answer the previous question. This complex matter will not be investigated further, since it goes beyond the scope of this thesis; refer to [63] for deeper insights.

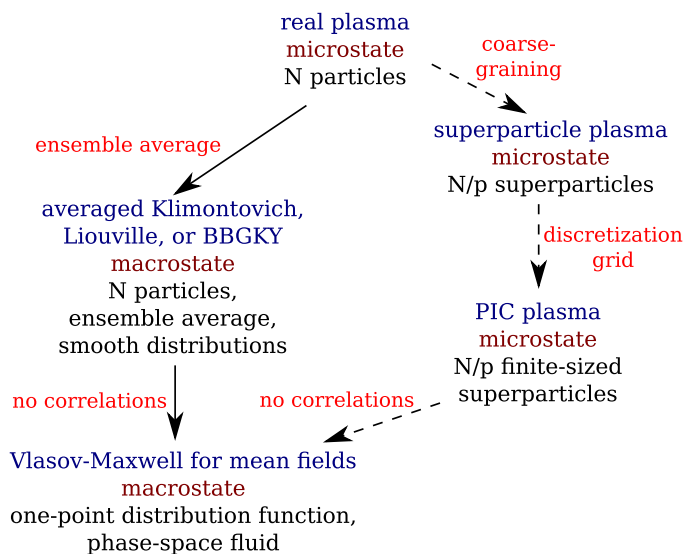


Figure 2.4: Different plasma models. Dashed arrows are transitions showing non-trivial effects. Reprinted from [63].

2.2.3 Problem definition and standard algorithms

Although it is not easy to determine the relations between Klimontovich, Vlasov and PIC descriptions, the problem solved by a PIC method is well-defined and will be presented in the following together with the basic numerical techniques used for its approximation.

Continuous PIC problem

From the foregoing it follows that every PIC method solves the Vlasov-Maxwell system without actually solving Vlasov equation. In fact, the problem of solving the complex partial differential equation that is Vlasov equation is reduced to the resolution of the system of ordinary differential equations that are the equations of motion for the macro-particles. The complete problem in continuous space-time reads as follows.

Let $\Omega \subset \mathbb{R}^3$ be the physical domain of the problem, $V \subset \mathbb{R}^3$ the volume enclosing the plasma. Choose a suitable shape function $S : \Omega \mapsto \mathbb{R}$. Define the initial conditions $\mathbf{r}_{p,a,0} \in \Omega, \mathbf{p}_{p,a,0} \in \mathbb{R}^3$ for all particles $p = 1, \dots, N_{p,a}$ of all species $a = 1, \dots, N$ and $\mathbf{E}_0, \mathbf{B}_0 : \Omega \mapsto \mathbb{R}^3$. Find for all $p = 1, \dots, N_{p,a}$ for all $a = 1, \dots, N$ the positions and momenta of the macro-particles $\mathbf{r}_{p,a}, \mathbf{p}_{p,a} : (0, T) \mapsto \mathbb{R}^3$ and for all $\mathbf{r} \in \Omega \subset \mathbb{R}^3$, for all $t \in (0, T)$ the electric and magnetic fields $\mathbf{E}, \mathbf{B} : \Omega \times (0, T) \mapsto \mathbb{R}^3$ such that:

$$\left\{ \begin{array}{ll}
\frac{d\mathbf{r}_{p,a}(t)}{dt} = \mathbf{v}_{p,a}(t) & \text{in } (0, T), \forall p = 1, \dots, N_{p,a}, \forall a = 1, \dots, N \\
\frac{d\mathbf{p}_{p,a}(t)}{dt} = \tilde{\mathbf{F}}_{p,a}(t) & \text{in } (0, T), \forall p = 1, \dots, N_{p,a}, \forall a = 1, \dots, N \\
\text{curl } \mathbf{E}(\mathbf{r}, t) = -\frac{1}{c} \frac{\partial \mathbf{B}(\mathbf{r}, t)}{\partial t} & \text{in } \Omega \times (0, T) \\
\text{curl } \mathbf{B}(\mathbf{r}, t) = \frac{4\pi}{c} \mathbf{J}(\mathbf{r}, t) + \frac{1}{c} \frac{\partial \mathbf{E}(\mathbf{r}, t)}{\partial t} & \text{in } \Omega \times (0, T) \\
\text{div } \mathbf{E}(\mathbf{r}, t) = 4\pi\rho(\mathbf{r}, t) & \text{in } \Omega \times (0, T) \\
\text{div } \mathbf{B}(\mathbf{r}, t) = 0 & \text{in } \Omega \times (0, T) \\
\mathbf{r}_{p,a}(0) = \mathbf{r}_{p,a,0} & \forall p = 1, \dots, N_{p,a}, \forall a = 1, \dots, N \\
\mathbf{p}_{p,a}(0) = \mathbf{p}_{p,a,0} & \forall p = 1, \dots, N_{p,a}, \forall a = 1, \dots, N \\
\mathbf{E}(\mathbf{r}, 0) = \mathbf{E}_0(\mathbf{r}) & \text{in } \Omega \\
\mathbf{B}(\mathbf{r}, 0) = \mathbf{B}_0(\mathbf{r}) & \text{in } \Omega,
\end{array} \right.$$

where

$$\begin{aligned}
\mathbf{p}_{p,a} &= m\gamma_{p,a}(t)\mathbf{v}_{p,a}(t), \quad \gamma_{p,a}(t) = \sqrt{1 + \frac{\mathbf{p}_{p,a}(t) \cdot \mathbf{p}_{p,a}(t)}{m_a^2 c^2}}, \\
\tilde{\mathbf{F}}_{p,a}(t) &= q_a \int_{\Omega} S(\mathbf{r} - \mathbf{r}_{p,a}(t)) \left[\mathbf{E}(\mathbf{r}, t) + \frac{\mathbf{v}_{p,a}(t)}{c} \times \mathbf{B}(\mathbf{r}, t) \right] d\mathbf{r}, \\
\mathbf{J}(\mathbf{r}, t) &= \frac{1}{V} \sum_{a=1}^N q_a \sum_{p=1}^{N_{p,a}} S(\mathbf{r} - \mathbf{r}_{p,a}(t)) \mathbf{v}_{p,a}(t), \quad \rho(\mathbf{r}, t) = \frac{1}{V} \sum_{a=1}^N q_a \sum_{p=1}^{N_{p,a}} S(\mathbf{r} - \mathbf{r}_{p,a}(t)).
\end{aligned}$$

The unknowns of the PIC problem are the positions and momenta $(\mathbf{r}_p(t), \mathbf{p}_p(t))$ of all macro-particles of all species, solutions of the equations of motion, and the electromagnetic field $(\mathbf{E}(\mathbf{r}, t), \mathbf{B}(\mathbf{r}, t))$, solution of the Maxwell equations. The two pairs of unknowns are related to each other through the current density and charge density fields $(\rho(\mathbf{r}, t), \mathbf{J}(\mathbf{r}, t))$, hence through the shape functions.

The total number of macro-particles N_p sampling the Vlasov fluid is assumed to be given. The choice of this parameter is non-trivial and problem-dependent. For computational

reasons N_p is almost always several orders of magnitude lower than the real number of particles, but it has to be large enough to limit the numerical noise. Each macro-particle of a given species has mass m and charge q with the same charge-to-mass ratio and the same mean charge density as in the real plasma.

As for all partial differential equations, a set of boundary conditions is needed as an additional constraint for well-posedness. Most of the time *periodic boundary conditions* are imposed both on the particles positions and the electromagnetic field. In Cartesian geometry, writing the domain as $\Omega = (x_1, x_2) \times (y_1, y_2) \times (z_1, z_2)$, it means that when a particle exits the domain, say in (x_1, y, z) , is re-injected from the opposite side, in this case (x_2, y, z) , with the same momentum. For what concerns the fields, say \mathbf{E} , the following must be satisfied:

$$\begin{cases} \mathbf{E}(x_1, y, z, t) = \mathbf{E}(x_2, y, z, t) & \forall y \in (y_1, y_2), \forall z \in (z_1, z_2), \forall t \in (0, T) \\ \mathbf{E}(x, y_1, z, t) = \mathbf{E}(x, y_2, z, t) & \forall x \in (x_1, x_2), \forall z \in (z_1, z_2), \forall t \in (0, T) \\ \mathbf{E}(x, y, z_1, t) = \mathbf{E}(x, y, z_2, t) & \forall x \in (x_1, x_2), \forall y \in (y_1, y_2), \forall t \in (0, T). \end{cases}$$

Unavoidable recirculation effects may arise imposing periodic boundary conditions. However, such conditions are easy to implement and in many situations the simulation box can be chosen big enough so that recirculation does not affect the results. Of course this may cause non-physical features to which one always has to pay attention. Another option is to choose *open boundary conditions*. In this case the particles leaving the domain are definitively lost, while the electromagnetic field satisfies far-field Maxwell equations on the domain boundaries. These conditions are more complex to implement, but lead to neat results for the electromagnetic field. On the other hand, they are not recommended for particles, since their losses may cause disruptive numerical instabilities because of the sudden disappearance of currents contributions.

As a final remark, it is possible to solve the problem without actually solving the equations for the divergences, since they only represent two constraint on the fields. If they are satisfied at the initial time, they remain correct for all times, given that the continuity equation is guaranteed as well for all times. This is easily verified, since:

$$\begin{aligned} \frac{\partial}{\partial t} \operatorname{div} \mathbf{B} &= \operatorname{div} \frac{\partial \mathbf{B}}{\partial t} = -c \operatorname{div} \operatorname{curl} \mathbf{E} \equiv 0 \\ \frac{\partial}{\partial t} (\operatorname{div} \mathbf{E} - 4\pi\rho) &= \operatorname{div} \frac{\partial \mathbf{E}}{\partial t} - 4\pi \frac{\partial \rho}{\partial t} = \operatorname{div}(c \operatorname{curl} \mathbf{B} - 4\pi\mathbf{J}) - 4\pi \frac{\partial \rho}{\partial t} = 0 \iff \frac{\partial \rho}{\partial t} + \operatorname{div} \mathbf{J} = 0. \end{aligned}$$

A PIC algorithm follows essentially from the discretization in space and time of the problem herein presented. To completely determine a PIC scheme one has to specify the algorithms used for the connection between particle-based and grid-based quantities, the solution of Maxwell equations and the solution of particles motion. In the following paragraphs the most commonly used algorithms in Cartesian geometry are presented. It is useful to

introduce a mesh of the Cartesian domain Ω whose nodes are identified by the triplets (i, j, k) , $i = 1, \dots, N_x$, $j = 1, \dots, N_y$, $k = 1, \dots, N_z$ with spacing Δx , Δy , Δz along each direction, respectively. In addition, the time interval is discretized in steps t^n , $n = 1, \dots, N_t$ with spacing Δt .

Particle-grid connection

Because the electromagnetic field is computed on a grid, some kind of regularization step is necessary for the coupling between grid-defined quantities and particle variables. This is a crucial step in a PIC method and concerns both the reconstruction of the charge and current densities on the grid from the particles velocities and positions and the computation of the Lorentz force acting on every particle from the grid defined fields \mathbf{E} and \mathbf{B} . Such connection is performed using the shape functions mentioned in paragraph 2.2.1. On Cartesian meshes the standard choice is to use B-splines, i.e. K -order polynomials of compact support, as shape functions.

Consider a 1D uniform grid with spacing h and points x_i , $i = 1, \dots, N_x$. The B-spline $S^K(x_i - x) = S^K(u)$ is the piece-wise polynomial with the following properties:

- $S^K \in C^{K-1}(\mathbb{R}) : S^K|_{[x_i, x_{i+1}]} \in \mathbb{P}^K([x_i, x_{i+1}])$
- compact support: $|u| \leq (K + 1)h/2$
- positivity: $S^K(u) \geq 0$
- normalization: $\int_{\mathbb{R}} S^K(u) du = 1$
- partition of unit: $h \sum_{i=1}^{N_x} S^K(x_i - x) = 1$
- parity: $S^K(-u) = S^K(u)$

The B-splines of order 0, 1, 2 and 3 are (see figure 2.5):

$$S^0(u) = \begin{cases} \frac{1}{h} & |u| < \frac{h}{2} \\ 0 & \text{else} \end{cases} \quad S^1(u) = \frac{1}{h} \begin{cases} 1 - \frac{|u|}{h} & |u| \leq h \\ 0 & \text{else} \end{cases}$$

$$S^2(u) = \frac{1}{h} \begin{cases} \frac{1}{2} \left(\frac{3}{2} - \frac{|u|}{h} \right)^2 & \frac{1}{2}h \leq |u| \leq \frac{3}{2}h \\ \frac{3}{4} - \left(\frac{|u|}{h} \right)^2 & |u| \leq \frac{1}{2}h \\ 0 & \text{else} \end{cases} \quad S^3(u) = \frac{1}{6h} \begin{cases} \left(2 - \frac{|u|}{h} \right)^3 & \frac{1}{2}h \leq |u| \leq 2h \\ 4 - 6 \left(\frac{|u|}{h} \right)^2 + 3 \left(\frac{|u|}{h} \right)^3 & 0 \leq |u| \leq h \\ 0 & \text{else} \end{cases}$$

The contribution (or weight) to a grid-based quantity on the i -th grid point of the particle




Figure 2.5: Splines of different order used as convolution kernels to couple particles with fields. Left: order 0, nearest grid point method. Middle: order 1, cloud-in-cell method. Right: order 2.

located in x_p is defined using the splines as:

$$B_i(x_p) = hS^K(x_i - x_p).$$

It follows that

$$\sum_{i=1}^N B_i(x_p) = 1.$$

Intuitively, the splines describe the shape of the density spatial profile of every macro-particle. If K -order splines in 1D are used, then every particle contributes to the total density only on the $K + 1$ grid nodes centred around the nearest node to the particle. Typically the order K of the spline is chosen between 1, 2 and 3. As an example, if $K = 2$ a particle is connected to 3 different grid points: the nearest to its position and the two adjacent ones. According to the 2-order splines, if the nearest node is x_j , the weight on x_j is given by $\frac{3}{4} - (\frac{x_p - x_j}{h})^2$, while those on $x_{j\pm 1}$ are given by $\frac{1}{2}(\frac{1}{2} \pm \frac{x_p - x_j}{h})^2$.

For a multidimensional Cartesian grid the particle contribution on the (i, j, k) node is given by the tensor product of one-dimensional splines along each coordinate direction:

$$\hat{B}_{i,j,k}(\mathbf{r}_p) = B_i(x_p)B_j(y_p)B_k(z_p)$$

with

$$\sum_{i,j,k} \hat{B}_{i,j,k}(\mathbf{r}_p) = 1.$$

Knowing the contribution of each particle on the grid points, the number density for each species on the grid is:

$$n_a(i, j, k) = \frac{1}{\Delta\Omega} \sum_{p=1}^{N_p} \hat{B}_{i,j,k}(\mathbf{r}_p)$$

where $\Delta\Omega = \Delta x \Delta y \Delta z$. The total charge density on the grid is:

$$\rho_{i,j,k} = \sum_{a=1}^N q_a n_a(i, j, k)$$

where q_a is the charge of each macro-particle of the species a . A possible algorithm for the computation of the current density field is given by

$$\mathbf{J} = \sum_{a=1}^N \mathbf{J}_a \quad \mathbf{J}_a = \frac{q_a}{\Delta\Omega} \sum_{p=1}^{N_p} \mathbf{v}_p \hat{B}_{i,j,k}(\mathbf{r}_p). \quad (2.4)$$

A similar procedure is used to compute the fields $\mathbf{E}(\mathbf{r}_p)$, $\mathbf{B}(\mathbf{r}_p)$ acting on each macro-particle from their values on the grid:

$$\mathbf{E}(\mathbf{r}_p) = \sum_{i,j,k} \mathbf{E}_{i,j,k} \hat{B}_{i,j,k}(\mathbf{r}_p), \quad \mathbf{B}(\mathbf{r}_p) = \sum_{i,j,k} \mathbf{B}_{i,j,k} \hat{B}_{i,j,k}(\mathbf{r}_p).$$

It can be proven that (see [64]), as long as the fields assignment to particle positions is formulated with the same splines used in the current density field, then (2.4) is an energy-conserving algorithm for current density deposition, meaning that is consistent with the particle energy balance given by equation (1.15). However, such expression for \mathbf{J} never ensures the continuity equation for the electric charge at the discrete level because of unavoidable discrepancies between the weighting on the grid of the charge and current density fields. This issue will be discussed further in section 4.2.

Maxwell solver

Numerical approximation of Maxwell equations is a well-studied problem common to many discipline, on which wide literature exist (see [65, 66, 67]). The numerical integration of Maxwell equations in PIC codes is often performed using a second order finite difference time domain (FDTD) solver on a Yee-lattice, first proposed by Yee in [68] (see also [69]). It is based on a leap-frog scheme in time and on a spatial discretization on a staggered grid to increase numerical accuracy. Second order accurate centred finite differences are used for both space and time derivatives. This scheme employs both integer and half-integer points in both space and time discretizations, defined as:

$$\begin{aligned} t^n &= n\Delta t & t^{n+1/2} &= (n + 1/2) \Delta t \\ x_i &= x_{min} + i\Delta x & x_{i+1/2} &= x_{min} + (i + 1/2) \Delta x \\ y_j &= y_{min} + j\Delta y & y_{j+1/2} &= y_{min} + (j + 1/2) \Delta y \\ z_k &= z_{min} + k\Delta z & z_{k+1/2} &= z_{min} + (k + 1/2) \Delta z. \end{aligned}$$

In other words, the staggered grid consists of two meshes shifted by half-cell along the three directions; an analogous consideration can be made on the time discretization. The electromagnetic field is approximated in such a way that the \mathbf{B} field is shifted in time by $\Delta t/2$ with respect to \mathbf{E} field. Thus the time evolution of \mathbf{E} from step n to step $n + 1$ depends on the values of the \mathbf{B} field at time step $n + 1/2$, while the time evolution of \mathbf{B} from step $n + 1/2$ to step $n + 3/2$ depends on the values of the \mathbf{E} field at time step $n + 1$,

and so on. In order to use second order centred finite differences for the discretization of the spatial derivatives, the fields components are arranged on a 3D grid as shown in figure 2.6.




Figure 2.6: Yee-lattice. Reproduced from Wikipedia.

Summarizing, \mathbf{E} and \mathbf{B} fields are computed at the time steps and on the grid points as follows:

$$\begin{aligned}\mathbf{E}(\mathbf{r}, t) &\approx \left(E_{x_{i+1/2, j, k}}^n \quad E_{y_{i, j+1/2, k}}^n \quad E_{z_{i, j, k+1/2}}^n \right) \\ \mathbf{B}(\mathbf{r}, t) &\approx \left(B_{x_{i, j+1/2, k+1/2}}^{n+1/2} \quad B_{y_{i+1/2, j, k+1/2}}^{n+1/2} \quad B_{z_{i+1/2, j+1/2, k}}^{n+1/2} \right).\end{aligned}$$

Consistently with this choice, the current density field and charge density field are collocated in space and time as follows:

$$\begin{aligned}\mathbf{J}(\mathbf{r}, t) &\approx \left(J_{x_{i+1/2, j, k}}^{n+1/2} \quad J_{y_{i, j+1/2, k}}^{n+1/2} \quad J_{z_{i, j, k+1/2}}^{n+1/2} \right), \\ \rho(\mathbf{r}, t) &\approx \rho_{i, j, k}^n\end{aligned}$$

As already mentioned, the computation of the charge density field may be avoided, since it is only necessary to solve the equations for the curls of \mathbf{E} and \mathbf{B} fields. The electromagnetic fields are explicitly computed solving the discretized versions of the Ampère-Maxwell and Faraday equations. A semidiscretization in time of these equations yields to:

$$\begin{aligned}\mathbf{B}^{n+1/2} &= \mathbf{B}^{n-1/2} - \Delta t(\text{curl } \mathbf{E})^n \\ \mathbf{E}^{n+1} &= \mathbf{E}^n + \Delta t \left((\text{curl } \mathbf{B})^{n+1/2} - 4\pi \mathbf{J}^{n+1/2} \right)\end{aligned}$$

where $(\text{curl } \mathbf{E})^n$ is computed on the same grid nodes of the corresponding \mathbf{B} component, while $(\text{curl } \mathbf{B})^{n+1/2}$ is computed on the same grid nodes of the corresponding \mathbf{E} component.

The complete scheme follows considering also space discretization applying second order centred finite differences for the approximation of the curls of the fields:

$$\begin{aligned}
(\text{curl } \mathbf{E})^n &= \begin{pmatrix} \frac{E_{z_{i,j+1,k+1/2}}^n - E_{z_{i,j,k+1/2}}^n}{\Delta y} - \frac{E_{y_{i,j+1/2,k+1}}^n - E_{y_{i,j+1/2,k}}^n}{\Delta z} \\ -\frac{E_{z_{i+1,j,k+1/2}}^n - E_{z_{i,j,k+1/2}}^n}{\Delta x} + \frac{E_{y_{i,j+1/2,k+1}}^n - E_{y_{i,j+1/2,k}}^n}{\Delta y} \\ \frac{E_{y_{i+1,j+1/2,k}}^n - E_{y_{i,j+1/2,k}}^n}{\Delta x} - \frac{E_{x_{i+1/2,j+1,k}}^n - E_{x_{i+1/2,j,k}}^n}{\Delta y} \end{pmatrix} \\
(\text{curl } \mathbf{B})^{n+1/2} &= \begin{pmatrix} \frac{B_{z_{i+1/2,j+1/2,k}}^{n+1/2} - B_{z_{i+1/2,j-1/2,k}}^{n+1/2}}{\Delta y} - \frac{B_{y_{i+1/2,j,k+1/2}}^{n+1/2} - B_{y_{i+1/2,j,k-1/2}}^{n+1/2}}{\Delta z} \\ -\frac{B_{z_{i+1/2,j+1/2,k}}^{n+1/2} - B_{z_{i-1/2,j+1/2,k}}^{n+1/2}}{\Delta x} + \frac{B_{y_{i+1/2,j,k+1/2}}^{n+1/2} - B_{y_{i+1/2,j,k-1/2}}^{n+1/2}}{\Delta y} \\ \frac{B_{y_{i+1/2,j,k+1/2}}^{n+1/2} - B_{y_{i-1/2,j,k+1/2}}^{n+1/2}}{\Delta x} - \frac{B_{x_{i,j+1/2,k+1/2}}^{n+1/2} - B_{x_{i,j-1/2,k+1/2}}^{n+1/2}}{\Delta y} \end{pmatrix}.
\end{aligned}$$

In order to guarantee numerical stability, the Courant condition between the time step and space grid sizes must be satisfied:

$$\frac{\Delta t}{\sqrt{\Delta x^2 + \Delta y^2 + \Delta z^2}} = \sigma < 1,$$

where σ is known as Courant factor.

Particle motion

The most used method for solving the motion of the macro-particles is a leap-frog algorithm in which position and velocity variables are discretized on a staggered time grid. The integration scheme for every particle is the following:

$$\mathbf{p}^{n+1/2} = \mathbf{p}^{n-1/2} + \Delta t \frac{q}{m} [\mathbf{E}^n + \mathbf{v}^n \times \mathbf{B}^n] \quad (2.5)$$

$$\mathbf{r}^{n+1} = \mathbf{r}^n + \Delta t \mathbf{v}^{n+1/2}, \quad (2.6)$$

where \mathbf{E}^n and \mathbf{B}^n are respectively the electric and magnetic field interpolated at the centre of the macroparticle located in \mathbf{r}^n at time $t^n = n\Delta t$. The \mathbf{B} field is only known at half-integer times, hence to get the value of \mathbf{B}^n , \mathbf{B} field must be advanced in two half-integer times. Momentum and velocity are only known at half-integer times as well, their expression at time n can be approximated by:

$$\begin{aligned}
\mathbf{v}^n &= \frac{\mathbf{p}^n}{\gamma^n} \\
\mathbf{p}^n &= \frac{\mathbf{p}^{n+1/2} + \mathbf{p}^{n-1/2}}{2}.
\end{aligned}$$

From the second $\mathbf{p}^{n+1/2}$ can be written in terms of \mathbf{p}^n and substituting in 2.5 the following implicit equation for \mathbf{p}^n is obtained:

$$\mathbf{p}^n = \mathbf{p}^{n-1/2} + \frac{\Delta t}{2} \frac{q}{m} \left[\mathbf{E}^n + \frac{\mathbf{p}^n}{\gamma^n} \times \mathbf{B}^n \right].$$

This equation can be solved explicitly using the Boris pusher algorithm as follows. Let

$$\begin{aligned} \beta &= \frac{q}{m} \frac{\Delta t}{2} \\ \mathbf{b} &= \beta \frac{\mathbf{B}^n}{\gamma^n} \\ \tilde{\mathbf{p}} &= \mathbf{p}^{n-1/2} + \beta \mathbf{E}^n \end{aligned}$$

hence

$$\mathbf{p}^n = \tilde{\mathbf{p}} + \mathbf{p}^n \times \mathbf{b}. \quad (2.7)$$

Noting that

$$\mathbf{p}^n \cdot \mathbf{p}^n = \tilde{\mathbf{p}} \cdot \tilde{\mathbf{p}} + O(\Delta t^2),$$

it is possible to approximately evaluate γ^n within the scheme approximation order as

$$\gamma^n = \sqrt{1 + \tilde{\mathbf{p}} \cdot \tilde{\mathbf{p}}}.$$

By right vector multiplying $\times \mathbf{b}$ equation (2.7) for the half advance of the particle momenta and after some computations:

$$\mathbf{p}^n = \frac{1}{1 + b^2} (\tilde{\mathbf{p}} + \tilde{\mathbf{p}} \times \mathbf{b} + \mathbf{b}(\tilde{\mathbf{p}} \cdot \mathbf{b})),$$

then the final solution is given by

$$\mathbf{p}^{n+1/2} = 2\mathbf{p}^n - \mathbf{p}^{n-1/2}.$$

From $\mathbf{p}^{n+1/2}$ it is possible to update the particle velocity

$$\mathbf{v}^{n+1/2} = \frac{\mathbf{p}^{n+1/2}}{\gamma^{n+1/2}}$$

and advance the particle position

$$\mathbf{r}^{n+1} = \mathbf{r}^n + \Delta t \mathbf{v}^{n+1/2}.$$

Overall this algorithm guarantees second order accuracy in time.

Main loop

Every PIC code consists of a main temporal loop, at each iteration of which the unknown variables are computed. After plasma and fields initialization, the positions and momenta of all particles along with the electromagnetic field on the grid nodes are computed for every time step. In a single iteration of the temporal loop the following steps are performed:

- evaluation of the current density field on the grid nodes from particles charges, positions and velocities;
- computation of the curls of the electromagnetic field;
- one step advance of the electromagnetic field using the current density and the curls;
- field assignment from the grid nodes to the particles positions;
- evaluation of the Lorentz force acting on each particle;
- one step evolution of particles momenta and positions.

As an example, figure 2.7 shows the routines called in the main loop of the code *piccante*, which implements the FDTD on a staggered lattice and Boris pusher algorithms with periodic boundary conditions (PBC).

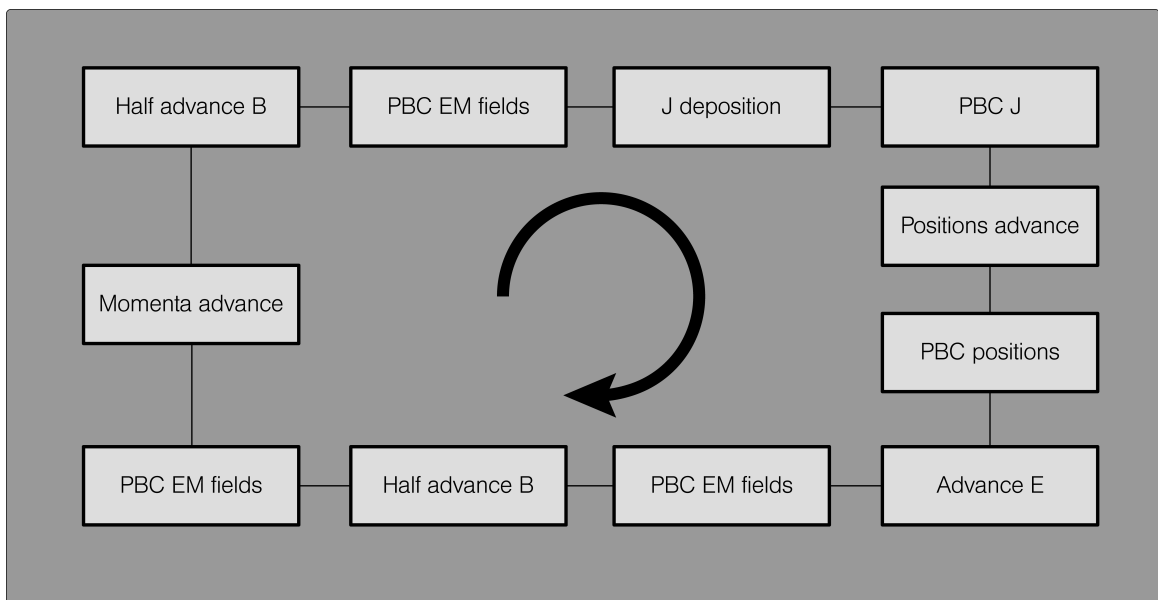


Figure 2.7: Main time cycle in a PIC code using standard algorithms.

2.2.4 Variations of the standard PIC method

In the previous sections the main features of PIC codes have been pointed out, along with the standard choices for the core algorithms. Here are summarized some other possible choices and variations of the basic PIC method.

The current deposition strategy based on the shape functions used in (2.4) is consistent with the single particle energy balance, but does not guarantee the continuity equation, i.e. charge conservation. A charge-conserving algorithm for the deposition of the current density vector from the particles to the grid can be obtained by forcing the continuity equation. This scheme, known as Esirkepov method [70], will be explained in detail in section 4.2. Another example of charge-conserving scheme is the Villasenor-Buneman method proposed in [71].

The second-order leap-frog integration scheme in space and time is easy to implement and efficient, however higher accuracy is gained only significantly increasing the space-time resolution. In some situations a high accuracy may be needed to prevent undesired numerical effects, such as poor energy conservation or phase velocity errors in long time simulations. In these cases, higher order integrations schemes, e.g. Candy-Rozmus [72] or Runge-Kutta schemes, can be exploited to have high accuracy with limited computational requirements. Refer to [64] for more details.

Besides the FDTD method, a multiplicity of Maxwell solvers exists, for example based on the finite element technique or on spectral methods. As a remark, explicit methods – above all the FDTD method – are the most widely employed in simulations of laser-plasma interaction, due to their efficiency. However, numerical artefacts may arise which can be suppressed using either a different method or some ad hoc strategies.

The δf method is a variant of the PIC scheme often used when the physics of the system remains close to a certain equilibrium, for example in tokamak plasmas or particle accelerators. The idea is to write the distribution function as $f = f_0 + \delta f$, where f_0 is the equilibrium configuration, and compute the perturbation δf with a PIC method [73]. This method significantly reduces the noise where $\delta f \ll f_0$. Another kind of particle algorithm was proposed in [74], trying to combine the best features of PIC and hydrodynamic codes. The idea is to push a relatively small (if compared to PIC) number of macro-particles that can directly interact with each other, split and merge.

The algorithms explained in the previous sections constitute the core of a PIC code for collisionless plasma simulation. However, collisional effects cannot be neglected in some situations, for instance at relatively low temperatures ($\ll 1KeV$) and high densities ($\gg 10^{21}cm^{-3}$). Moreover, especially in high intensity laser-plasma interaction investigations, electron ionization by a high intensity laser and through collisions along with nonlinear

quantum electrodynamics (QED) emission processes can play an important role in the physics of the system. For these reasons, modern PIC codes implement new packages that enable additional physical features, such as binary collisions, ionisation and QED effects. More details are presented in [75].

2.3 Comparisons

Many important differences exist between grid-based methods and the PIC scheme, conditioning the choice of the method depending on the addressed problem.

Grid-based methods guarantee a fine and uniform resolution of the phase space, even in those regions where the distribution function takes very small values. The price for a high fidelity solution is the greater numerical complexity and higher computational cost. Moreover, as opposed to PIC codes, they are not affected by numerical noise. Indeed, Vlasov codes are too expensive to be exploited in full dimensionality simulations, for which PIC codes remain basically the only option. Although not affected by numerical noise, grid methods suffer from artificial diffusion and the inability to resolve the appearance of small scales during the computation.

The PIC scheme is by far the most used for the Vlasov-Maxwell system approximated solution, especially in the laser-plasma community. Its strengths are robustness, flexibility, high scalability and, above all, great simplicity. The main drawback is the inherent numerical noise due to the sampling of the distribution function through a limited number of macro-particles. Important numerical noise may prevent from accurately simulating some phenomena, especially those involving small perturbations or those deeply affected by the tails of the distribution function. Furthermore, the limited number of macro-particles implies a minimum value of the density that can be accurately resolved. However, current PIC codes include additional physical features, greatly broadening the spectrum of simulable phenomena.

In conclusion, both Vlasov and PIC codes represent an invaluable numerical tool for collisionless plasma physics investigations. Vlasov codes are useful when high resolution information has to be extracted from the distribution function, whereas PIC codes allow high dimensionality simulation of a wide range of physical phenomena as well as the performance of extensive numerical campaigns. As the computational capabilities of present-day supercomputers increase, grid-based multidimensional simulations are becoming possible, making Vlasov codes an alternative to PIC codes, especially for benchmarking. Vlasov codes can compete with PIC codes in terms of computational cost in very few cases, when high resolution in reduce dimensionality is required [76, 77].

2.4 Motivations and goals of this thesis work

At this point I have presented all the elements necessary to understand the motivations and goals of my thesis work, but first let me overview what has been discussed thus far. Chapter 1 was devoted to the main mathematical aspects underlying the Vlasov-Maxwell problem and the discussion of some physical phenomena that can be described by those equations. Special attention has been given to the physics of laser-plasma interaction and, in particular, to laser-driven ion acceleration and its enhancement through special nanostructured target materials. Once depicted the problem under consideration, the tools used for its resolution have been presented in chapter 2, focused on the numerical methods for the approximation of Vlasov-Maxwell system, with particular care towards particle-in-cell codes.

The motivations of this work come from the ongoing investigations in the field of laser-driven ion acceleration, an appealing topic both for its applicative potential and the related studies on new regimes of radiation-matter interaction. The system under examination consists of a high intensity laser pulse irradiated onto a target that rapidly becomes plasma. The resulting interaction process is complex and strongly nonlinear. Moreover, additional complexities arise when choosing special targets, for instance nanostructured materials in multi-layer configurations. The most suitable description for this kind of physical situations is given by the kinetic theory. Collisionless approximation, i.e. Vlasov-Maxwell system, is adopted most of the times, because of the high temperatures (\gtrsim keV) at play and small time scales (~ 100 fs) and length scales (~ 100 nm \div few μ m) of interest. The exploitation of reduced models such those exposed at the beginning of section 2.1 is not worthwhile: the whole Vlasov-Maxwell system has to be solved not to lose important features of the phenomena, especially when considering near-critical, nanostructured targets with complex spatial non-uniformities. Under such conditions, multidimensional PIC codes represent the proper numerical tool to perform extensive parametric scans rather than Vlasov codes. It is worth mentioning that code optimization and, more in general, the exploitation of high performance computing paradigms is essential to be able to carry out such simulations. It is within this framework that I developed my thesis work, which consists of two main, parallel goals. Briefly they are: 1) exploitation of an existing particle-in-cell code for the simulation of laser-driven ion acceleration experiments with nanostructured targets and 2) implementation of new features and additional optimizations into the same code.

First, I exploited the particle-in-cell code *piccante* [78] to set and analyse a batch of two-dimensional and (very expensive) three-dimensional numerical simulations for laser-driven ion acceleration investigations. The aim of this part of the work is to gain more insights on the interaction and acceleration processes, employing different nanostructured multi-layer targets. I considered both idealized and more realistic models for the nanostructured

layer and let vary some of their characteristics. I chose simulation parameters comparable to those used in recent experiments in order to be able to compare simulations output with experimental results. Moreover, not many three-dimensional numerical studies with nanostructured targets in the near-critical regime can be found in literature. Therefore the results of these simulations are available for comparisons with other past and possible future experiments. A detailed discussion of these activities can be found in chapter 3.

Second, I worked on the improvement of the code *piccante*, implementing two additional features: shared memory parallelism (OpenMP) and a charge-conserving current deposition method. Part of the work concerned the partial parallelization of the code, on which a distributed memory paradigm was already implemented (MPI). Numerical simulations of laser-plasma interaction phenomena of actual interest require a lot of computational resources, thus are performed on supercomputers. To fully exploit these architectures and to afford multidimensional simulations on such machines, code parallelization is crucial. Depending on the architecture, optimum performances may be achieved exploiting both distributed-memory and shared-memory (MPI and OpenMP) models. For this reason I added the OpenMP paradigm to the code and compared the performances with MPI on different machines and with different simulation setups. Moreover, *piccante* only implemented an energy-conserving current deposition scheme, according to which the current density field is straightly scattered onto the grid through shape functions. Hence, it was important to understand whether the lack of exact charge conservation had negative consequences on the simulation of laser-driven ion acceleration experiments or not. For this reason I implemented the Esirkepov method (see section 2.2.4) into *piccante* and compared the results with those obtained using the standard procedure. This part of the work is presented in chapter 4.

Lastly, in chapter 5 I will highlight the conclusions of this work and its future perspectives.

Chapter 3

Numerical simulations of laser-plasma interaction with nanostructured targets

This chapter is devoted to the numerical investigation of the interaction between a high-intensity laser pulse and nanostructured plasmas. The models used to describe and simulate the foam-attached targets for ion acceleration described in 1.5.2 are explained in section 3.1. The objectives of the numerical campaigns along with the tools used to perform them are presented in section 3.2. The code used to perform all the simulations is presented in section 3.3. In section 3.4, the numerical results of three and two-dimensional particle-in-cell simulations are discussed. Lastly, section 3.5 illustrates some of the issues related to the computational cost of such simulations that naturally lead to the work presented in chapter 4.

3.1 Modelling nanostructured targets for laser-driven ion acceleration

In the framework of laser-driven ion acceleration, target characteristics play a crucial role in the physics of the interaction (see chapter 1). The choice of proper non-conventional targets may lead to different, complex dynamics and eventually to the enhancement of the acceleration process. For this reason, particular attention was given to novel target concepts and modelling. Inspired by recent experimental activities [32, 33], numerical simulations focused on special multi-layer targets with particular density, thickness, morphology and structure. Numerical simulations have been designed taking into account the experimental concepts discussed in section 1.5.2, so to be able to compare numerical results




Figure 3.1: Simplified scheme of a numerical multi-layer target: foam (dark blue), solid foil (red) and contaminant layer (light blue). Typical values of the electron density and thickness are indicated.

with experimental results.

Often, foam-attached targets for laser-driven ion acceleration are numerically modelled as three layer targets (see figure 3.1) [31]. The first layer (dark-blue) is located on the front side with respect to the laser propagation direction, i.e. on the illuminated side. It is a near-critical density layer of few-microns thickness representative of the foam, whereby the target is non-conventional, since it is added to enhance the acceleration process with respect to the case of simple solid target, as explained in section 1.5.2. The first layer is attached to a second layer (red) which consists in a solid μm -thick foil as those conventionally used in TNSA experiments. A third layer (light blue) is located on the rear side of the second one. It is a thin (tens of nanometres thick) impurity layer made of protons. This is to represent those contaminant protons present on the surface of the solid that are in a favourable condition to be accelerated because of both their optimal position and high charge-to-mass ratio that makes them move faster than all other ions (see section 1.5). Simple solid targets, i.e. without the foam-attached layer on the front side, are considered as a reference case for comparisons.

Different models for the foam layer are possible. A rough approximation is to consider the foam as a uniform layer at near-critical density. This implies the complete disregard of the effects related to the inhomogeneities of the foam; it is a reasonable approximation if the scale lengths of the foam microscopic structures are smaller than the laser wavelength and the characteristic scale lengths of the related plasma dynamics. To partially take into account the effects of inhomogeneities, it is possible to add a constant density gradient so that the resulting layer has a ramp on its front side. A different option is to model the foam as made of nanospheres regularly or randomly arranged in space. This approach has the advantage of being closer to reality than using a uniform layer. Regularity is a fictitious

feature in relation to the usual growth processes for foam manufacturing and may introduce artificial effects; however, it is interesting to investigate its consequences. This model has the advantage of taking into account the presence of small, dense clusters. However, it does not reproduce properly the aggregation of such clusters on larger scales.

A more realistic foam description is obtained by introducing a specific mathematical model for the material growth, such as the *diffusion-limited aggregation model* (DLA) proposed by Witten and Sander in 1981 [79]. This model is used for the description and simulation of the process where particles irreversibly combine together to form aggregates on scales larger than their characteristic dimension, when diffusion is the main transport mechanism in the system. The algorithm that simulates such process is the following. A seed particle is located at the origin of a lattice. A second particle is allowed to randomly walk (due to Brownian motion, i.e. diffusion) on the lattice from a random site located far away from the origin until it visits a lattice point adjacent to the seed. Then the particle is halted and a two-particle cluster is produced. A third particle is injected in the lattice from a new random location and randomly walks until it becomes part of the already formed cluster. This procedure is iterated for as many particles as desired. The same algorithm is used for the generation of a numerical foam where the particles are spheres organized in space according to the DLA model, representative of the nanoparticles, i.e. the building blocks of the aggregates constituting the foam (see paragraph 1.5.2). The morphology obtained with this method displays tree-like macro-structures, similar to those exhibited by the foams grown experimentally by pulsed laser deposition. An example of such foams is shown in 3.2 (compare with 1.12). Different structures and filling factors

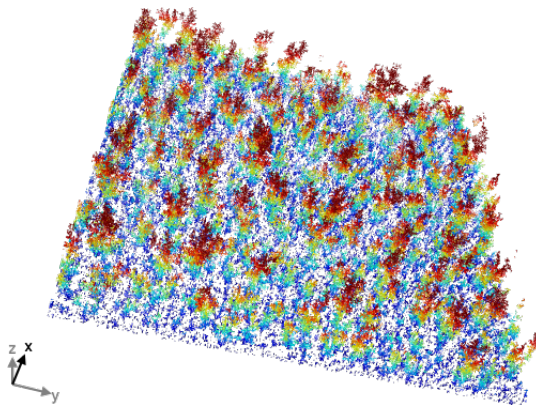


Figure 3.2: Graphical representation of a 3D DLA numerical foam.

can be obtained by properly tuning some parameters of the algorithm, such as the allowed initial positions and the probability of moving in each direction. For instance, a higher probability of moving towards the target leads to a more "compact" foam with higher filling factor. On the contrary, a higher probability of moving away from the target leads to an "emptier" structure with smaller filling factor. The same results are obtained by

respectively decreasing and increasing the initial distance from the target. To avoid too much inhomogeneities the foam can be grown up to a thickness bigger than desired and then cut.

In conclusion, the DLA model describes specific growth regimes and seizes the main features of the foam-attached targets production process. The model could be extended to take into account other physical aspects to simulate other realistic materials.

Only few numerical works on laser-plasma interaction with nanostructured targets can be found in literature, especially with foam-attached targets. In these articles foams are modelled either as uniform layers [80, 31, 32] or with the DLA method [33].

3.2 Goals and methods of the numerical campaign

Both three-dimensional and two-dimensional numerical simulations have been performed with the open source particle-in-cell code *piccante*, whose main features are described in section 3.3, to explore the physics of the interaction with different target configurations, namely to investigate the role of the foam structure on the process. Different multi-layer foam-attached targets have been modelled, changing the characteristics of the foam density distribution.

In every two-dimensional simulation a charged particle actually represents an infinite charged wire. For inherently 3D processes (such as TNSA) this often leads to the unphysical alteration of certain quantities, hence two-dimensional simulations can only provide qualitative information. For example when simulating TNSA-like processes in two dimensions, an enhancement in the energy of the accelerated ions occurs. Moreover, some physical features cannot be adequately modelled in two dimensions. A typical example is given by the pulse polarization: the electric field of a S polarized pulse lays outside the simulated plane, hence it cannot actually move the particles since they are infinite in that direction. As another example, it is very difficult to simulate high density ($\sim 100 n_c$), connected structures with low ($\sim 1 n_c$) average density. To collect quantitative estimations, three-dimensional simulations must be performed. Nevertheless, extensive parametric numerical studies cannot be carried out in three dimensions because of their extremely high computational cost. For this reason both 3D and 2D simulations have been performed, with different goals. The set of 3D simulations presented in subsection 3.4.1 aims to quantitatively analyse the acceleration performances and compare the results with available experimental data using quite realistic foam models, i.e. DLA numerical foams. It also allows to analyse the modelling of realistic properties related to the foam characteristics, such as density and thickness inhomogeneities. Besides, this numerical campaign also has the objective of establishing the computational burden and feasibility of such complex sim-

ulations so to be able to properly gauge the cost of possible future investigations. The set of 2D simulations in 3.4.2 constitutes a qualitative scan of more idealized foams, namely nanostructured foams made by dense spheres regularly arranged in space. This kind of simulations allow a broader analysis because of the simplifications lying both in the reduced dimensionality and the chosen foam model.

Three-dimensional simulations were carried out on Intel cluster Marconi, while two-dimensional simulations on Fermi supercomputer, both hosted at CINECA, the largest Italian computing centre located in Bologna.

Marconi is the new CINECA's Tier-0 machine, open to production since the 4th of July, 2016. It is a Lenovo NeXtScale system based on Intel technology, at present day equipped with 1512 compute nodes, each with 2 18-core Intel(R) Xeon(R) E5-2600 v4 processors running at 2.30 GHz and 128 GB of RAM. It will be gradually upgraded with the next-generation of the Intel Xeon Phi product family (Knights Landing) reaching a total computational power of about 20 PFlops.

Fermi was a Tier-0 machine for scientific research with system peak performance of 2.1 PFlops. It was a IBM Blue Gene/Q system configured with 10240 PowerA2 sockets running at 16 GHz, with 16 cores each, with 16GB of RAM per core. It has been switched off on July 15, 2016. Roughly, every CPUhour on Marconi corresponds to about 7 CPUhours on Fermi.

Simulations setup and results are presented in section 3.4.

3.3 *piccante*: a particle-in-cell code

piccante is a massively parallel, fully relativistic, electromagnetic, 3D particle-in-cell code, released under a GPLv3 license [78]. It has been developed by A. Sgattoni (CNR-INO), L. Fedeli (University of Pisa and CNR-INO) and S. Sinigardi (University of Bologna). The code is particularly suited for laser-plasma interaction investigations. It is written in C++ and is designed to run on a wide range of machines, from laptops to supercomputers. In the past few years, *piccante* has been exploited to investigate several laser-plasma interaction scenarios, such as laser-driven ion acceleration with innovative targets and high field plasmonics in solid targets [81, 82, 83, 84]. Here the key features of *piccante* are reported.

Simulation setup A user-friendly input file allows the configuration of the simulation setup. Simulations in 1D3V, 2D3V and 3D3V Cartesian geometries are allowed. An arbitrary number of plasma species and laser pulses can be placed inside the simulation box with different polarizations (P, S, C) and angles of incidence. The plasma can be made

by electrons, positrons or ions. Each species can be initialized with different boostable temperature distributions (Maxwellian, Waterbag and Supergaussian) and spatial density profiles, e.g. uniform, uniform with a ramp, collection of spheres, etc. Moreover, numerical DLA foams can be generated by a special routine that implements the DLA algorithm and writes in a binary file the coordinates of the spheres constituting the foam. *piccante* reads this file and consequently initializes the numerical plasma. The laser pulses can be Gaussian or a plane wave. The code is parallelized with MPI slicing the simulation box in sub-domains, each assigned to a different task. *piccante* has been tested on up to 32768 parallel processes and it has proven to be well scalable [85]. It is possible to stretch the simulation box in order to save memory while keeping high resolution where it is required. To save computational resources, simulations can be performed using a moving window, i.e. translating the simulation box with a chosen frequency. As an additional physics feature, radiation friction losses can be enabled (relevant for very high intensity $I > 10^{22} \text{ W/cm}^2$ laser-matter interaction). Flexible output is provided along with the tools necessary for its analysis. Heavyweight output files can contain charge density, current density, electromagnetic field and species phase space and can be limited to a specific region of the box. Lightweight output files contain synthetic information on the time evolution of some quantities (species total energy, E_x^2 , B_x^2 ,...) and species-related quantities (minimum and maximum momenta along each direction, Lorentz factor,...).

Algorithms *piccante* implements standard PIC algorithms that guarantee an overall second order accuracy both in space and time. The particle-grid connection is performed using second order polynomial shape functions, including the computation of the current density field. Hence, the current deposition strategy is energy-conserving. A FDTD (Finite Difference Time Domain) on a Yee-lattice is used for the integration of Maxwell equations. Periodic or open boundary conditions are enforced on the fields. Particles motion is solved using a leap-frog scheme and the Boris pusher algorithm with periodic boundary conditions on the positions. All these algorithms have already been discussed in detail in section 2.2.3.

Normalized units Normalized quantities are used in *piccante*. A physical characteristic length ℓ is set as the unit of length. When a laser pulse is present a common choice is to set $\ell = \lambda$, the laser wavelength. Time is normalized in units of ℓ/c , where c is the speed of light. This implies fixing the speed of light $c = 1$, thus space and time have the same units. Positions, velocities and momenta are normalized according to the following formulas (normalized quantities are indicated with *):

$$\mathbf{r}^* = \frac{\mathbf{r}}{\ell}, \quad \mathbf{p}^* = \frac{\mathbf{p}}{mc}, \quad \mathbf{v}^* = \frac{\mathbf{v}}{c} = \frac{\mathbf{p}^*}{\gamma}, \quad \gamma = \sqrt{1 + \mathbf{p}^* \cdot \mathbf{p}^*} = \frac{1}{\sqrt{1 - \mathbf{v}^* \cdot \mathbf{v}^*}}.$$

Charges are expressed in units of the elementary charge e , masses are expressed in units of the electron mass m_e . Therefore the electromagnetic field is normalized as follows:

$$\mathbf{E}^* = \frac{e \ell}{m_e c^2} \mathbf{E}, \quad \mathbf{B}^* = \frac{e \ell}{m_e c^2} \mathbf{B}.$$

The current density vector is expressed as:

$$\mathbf{J}^* = \frac{e \ell^2}{m_e c^3 \pi} \mathbf{J}.$$

Hence normalized Maxwell equations (in Gaussian CGS units) and particles equation of motion are (dropping the superscript *):

$$\left\{ \begin{array}{l} \text{curl } \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \\ \text{curl } \mathbf{B} = 4\pi^2 \mathbf{J} + \frac{\partial \mathbf{E}}{\partial t} \end{array} \right. \quad \left\{ \begin{array}{l} \frac{\partial \mathbf{r}}{\partial t} = \frac{\mathbf{p}}{\gamma} \\ \frac{\partial \mathbf{p}}{\partial t} = \frac{q}{m} (\mathbf{E} + \mathbf{v} \times \mathbf{B}). \end{array} \right.$$

When a laser pulse is present, the charge density field for each specie is often expressed in units of the critical density:

$$\rho^* = \frac{\rho}{en_c}.$$

3.4 Results

3.4.1 Three-dimensional PIC simulations

Real nanostructured foams described in section 1.5.2 are characterized by strong non-uniformities resulting from the production process. As a consequence, different outcomes (e.g. different cutoff proton energies) may be observed depending on properties of the illuminated region. For this reason, four large scale three-dimensional simulations have been carried out, varying the characteristics of the region hit by the laser pulse.

The simulated targets are:

- bare solid foil,
- uniform near-critical foam attached to a solid foil,
- nanostructured near-critical foam with a hollow in the illuminated region, attached to a solid foil ("hollow foam"),
- nanostructured near-critical foam with a hill in the illuminated region, attached to a solid foil ("hill foam").

Each target is irradiated by a high-intensity P polarized Gaussian laser pulse with $a_0 = 18$ at normal incidence, i.e. the laser pulse impinges perpendicularly on the target. This latter feature was settled for the sake of simplicity, since a zero angle of incidence is rarely chosen in experiments because of safety reasons (laser backscattering may damage the equipment). Assuming a wavelength of $0.8 \mu\text{m}$, the peak intensity is $I = 7 \times 10^{20} \text{ W/cm}^2$ ($a_0 = \sqrt{2e^2\lambda^2 I / \pi m_e^2 c^5}$). The laser pulse has a \cos^2 -function temporal profile with full-width-half-maximum of the field equal to 41 fs, equivalent to a full-width-half-maximum of the intensity of 30 fs. The laser focal spot size is $2.4 \mu\text{m}$.

The solid foil is a uniform layer with $0.8 \mu\text{m}$ thickness, made of electrons and ions with $Z/A = 27/11$, representative of Al^{11+} ions (with atomic number $Z=11$ and mass number $A=27$, not completely ionized). Its electron density is $50n_c$, sampled with 50 electrons per cell. The ion distribution is sampled with 8 particles per cell, much less than for the electron population. The reason for this choice is that the dynamics of this particular ion population is not relevant in such short time scales at this intensities. Thus it is crucial to solve only the electron density (with a high enough number of electrons per cell), but not the ion density. A thin (40 nm-thick) and low density ($5n_c$) contaminant layer is added on the rear side of the solid foil. It is made of protons – the accelerated species – and electrons, sampling both populations with 50 particles per cell. Here a finer sampling is required because statistical information on the contaminant ions, e.g. energy spectrum, is extracted.

All foams are made of electrons and ions whose charge-to-mass ratio is 2 (as for C^{6+} ions with $Z=6$, $A=12$); their electron density (mean if nanostructured) is $2n_c$. The uniform foam is a $4 \mu\text{m}$ -thick layer sampled with 4 electrons and 1 ion per cell, respectively. Figure 3.3 provides a graphical representation of the 3D nanostructured numerical foam targets. They are simulated as a collection of dense nanospheres arranged in space according to the diffusion-limited aggregation model as explained in subsection 3.1. The radius of the nanospheres is 40 nm, while their density is $50n_c$, sampled with 50 electrons and 9 ions per cell, respectively. The clusters are grown up to a maximum height of $9.6 \mu\text{m}$ and then cut at $4 \mu\text{m}$ to avoid excessive inhomogeneities. The resulting filling factor is about 0.04%, leading to a mean electronic density of $2n_c$. To obtain the hollow (hill) in the midst of the foam it is necessary to tune some parameters in the routine used to generate the foam itself. In particular, in the growth process the only allowed initial y and z coordinates for the spheres are defined outside (inside) a circle centred in the middle of the foam of radius $3.6 \mu\text{m}$ (refer to figure 3.3 to understand the coordinate system). Furthermore the initial heights (x coordinates) to obtain the hole have to be higher than to obtain the hill so that the particles have enough time to diffuse, hence fill the near-boundary region. Also the probability of moving towards regions with smaller x coordinate is higher for the hill case than for the hole case. The sizes of the hollow and the hill of the nanostructured foams are larger than the laser focal spot.

The simulation parameters are partially similar to those used in the experiments in [33]. A comparison between experimental and numerical results can give important quantitative insights, but a complete adherence between the two is not expected. For instance, the laser angle of incidence is different in the two situations.




Figure 3.3: Numerical foams obtained with the DLA method, with the hollow (left) and with the hill (right). Particles' colour follows the x coordinates.

The simulation time is 160 fs. All species of all plasmas are initialized with a Maxwellian distribution with temperature around 1 eV (10^4 K).

Table 3.1 summarizes the parameters common to all simulations. Figures 3.4, 3.5, 3.6 and 3.7 display a 3D snapshot for every simulation where electron density of the target and electromagnetic energy density are shown in a portion of the simulation box. Figures 3.8, 3.9, 3.10 and 3.11 show some graphical results of the numerical campaign. Each one is dedicated to a single case. Electron density and electromagnetic energy density sections on the symmetry plane ($z = 0$) are represented for three time steps on the left and right columns respectively. These figures show that the target is never seriously broken through, which means that only a small amount of the laser electromagnetic energy passes through the target, while in a large part is absorbed or reflected.

At the initial time the laser pulse begins to travel towards the target along the x direction at the speed of light. After few femtoseconds the pulse starts interacting with the foam, if present. Being the foam near-critical, the laser pulse can propagate through it and reach the solid substrate, where it is reflected backwards. It is worth noting that the foam layer is transparent to the laser pulse, due to relativistic effects (see section 1.4). The interaction with the solid layer begins around time $t = 40$ fs, when its electronic component starts to gain energy.

The laser pulse directly interacts only with the electron populations while, because of their much higher inertia to motion, ions dynamics is a consequence of electrons dynamics, i.e.

3D Simulation Parameters	
Box size	$61\lambda \times 80\lambda \times 80\lambda$
Resolution (points per λ)	$[44, 40, 40]$
Simulation time	60
Solid foil	
Density	$50n_c$
Thickness	1λ
Electrons per cell	50
Ions per cell	8
Contaminant layer	
Density	$5n_c$
Thickness	0.05λ
Electrons per cell	50
Protons per cell	50
Laser	
Type	Gaussian
Incidence	0°
Duration FWHM	$15.4\lambda/c$
Initial peak position	$[-16\lambda, 0, 0]$
Waist	3λ
a_0	18
Foams	
Thickness	5λ
Mean density	$2n_c$

Table 3.1: Summary of three-dimensional simulations setup.




Figure 3.4: Snapshot of the 3D simulation with the bare solid target ($t = 80$ fs).

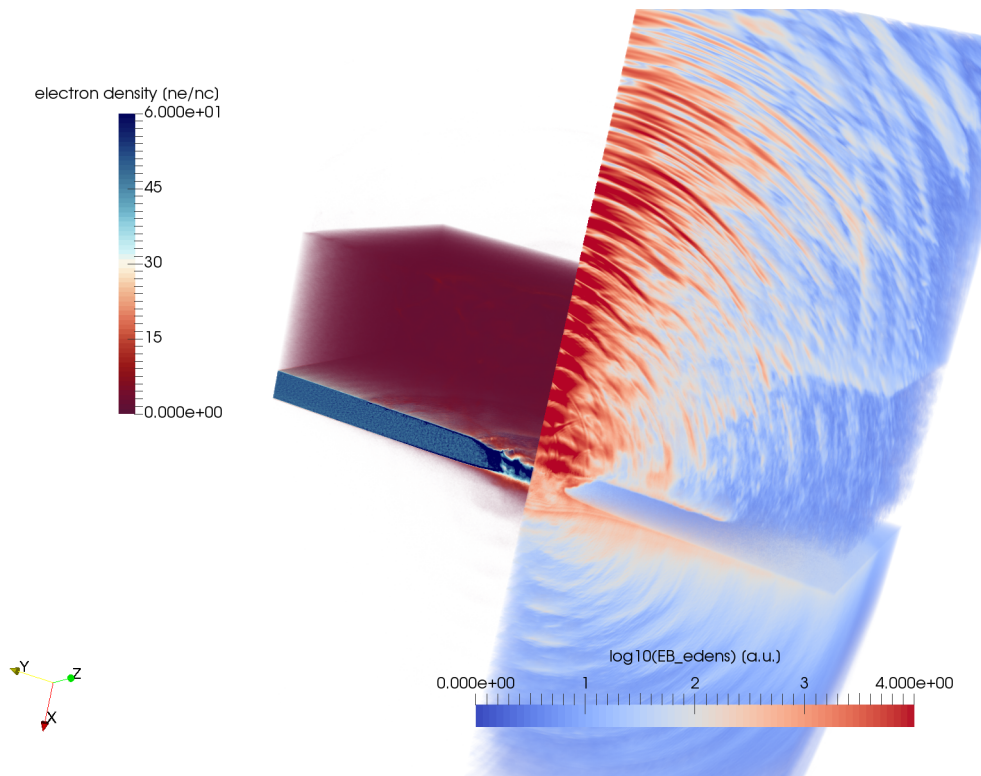


Figure 3.5: Snapshot of the 3D simulation with the uniform foam-attached target ($t = 80$ fs).




Figure 3.6: Snapshot of the 3D simulation with the hollow foam-attached target ($t = 80$ fs).

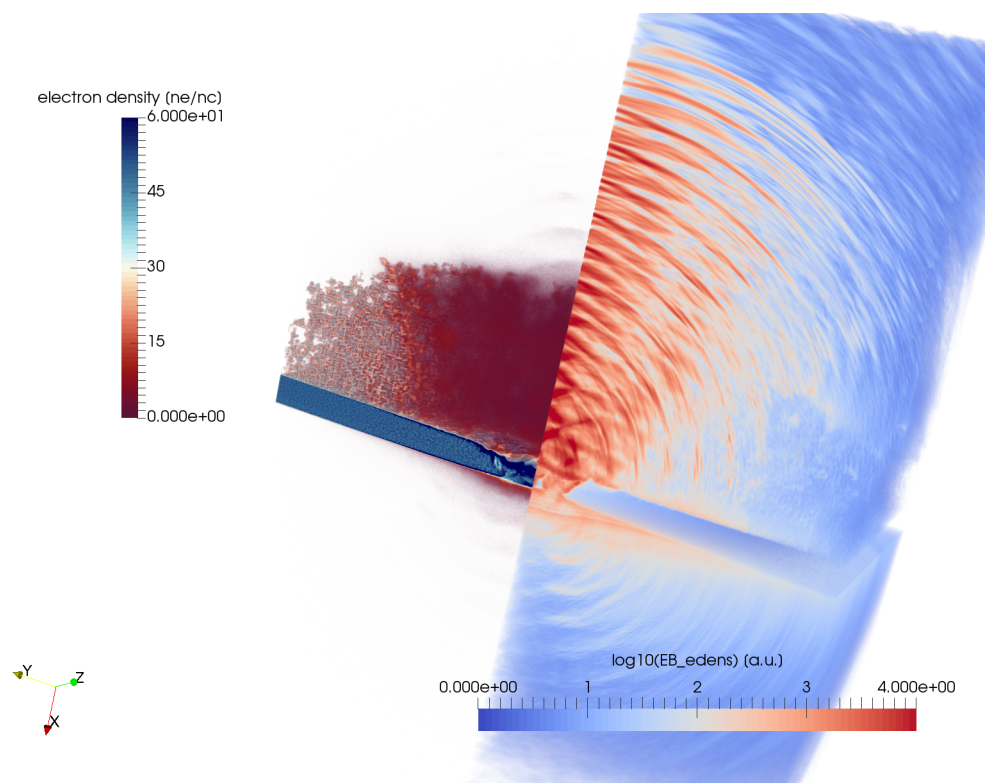


Figure 3.7: Snapshot of the 3D simulation with the hill foam-attached target ($t = 80$ fs).

NO FOAM




Figure 3.8: Electron density and electromagnetic energy density at plane $z = 0$.

UNIFORM FOAM

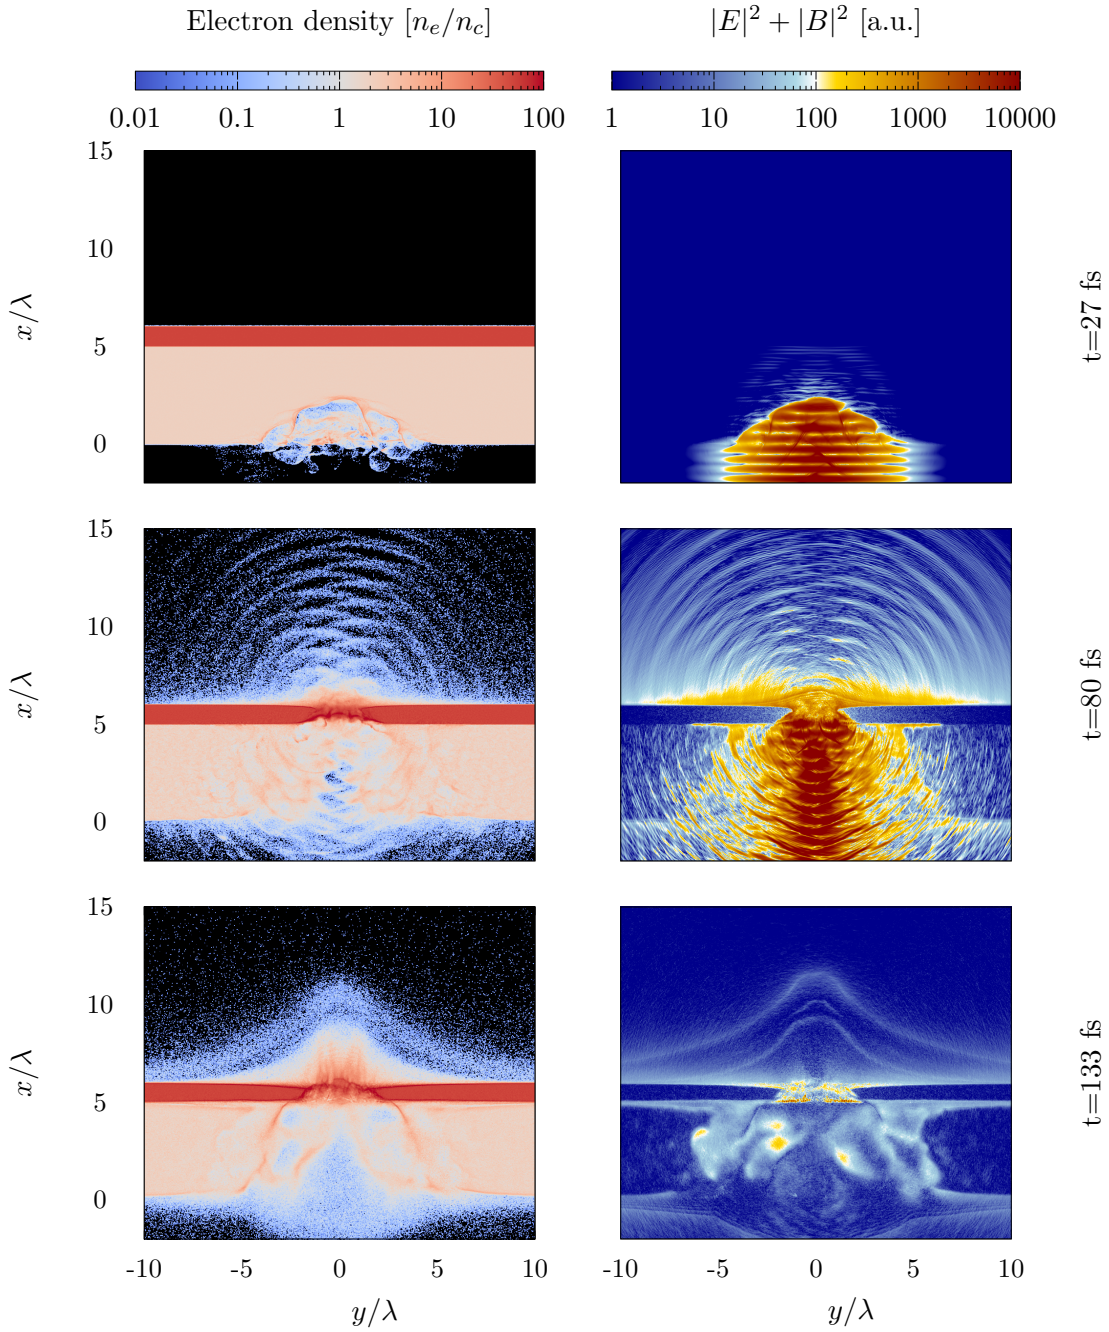


Figure 3.9: Electron density and electromagnetic energy density at plane $z = 0$.

HOLE FOAM

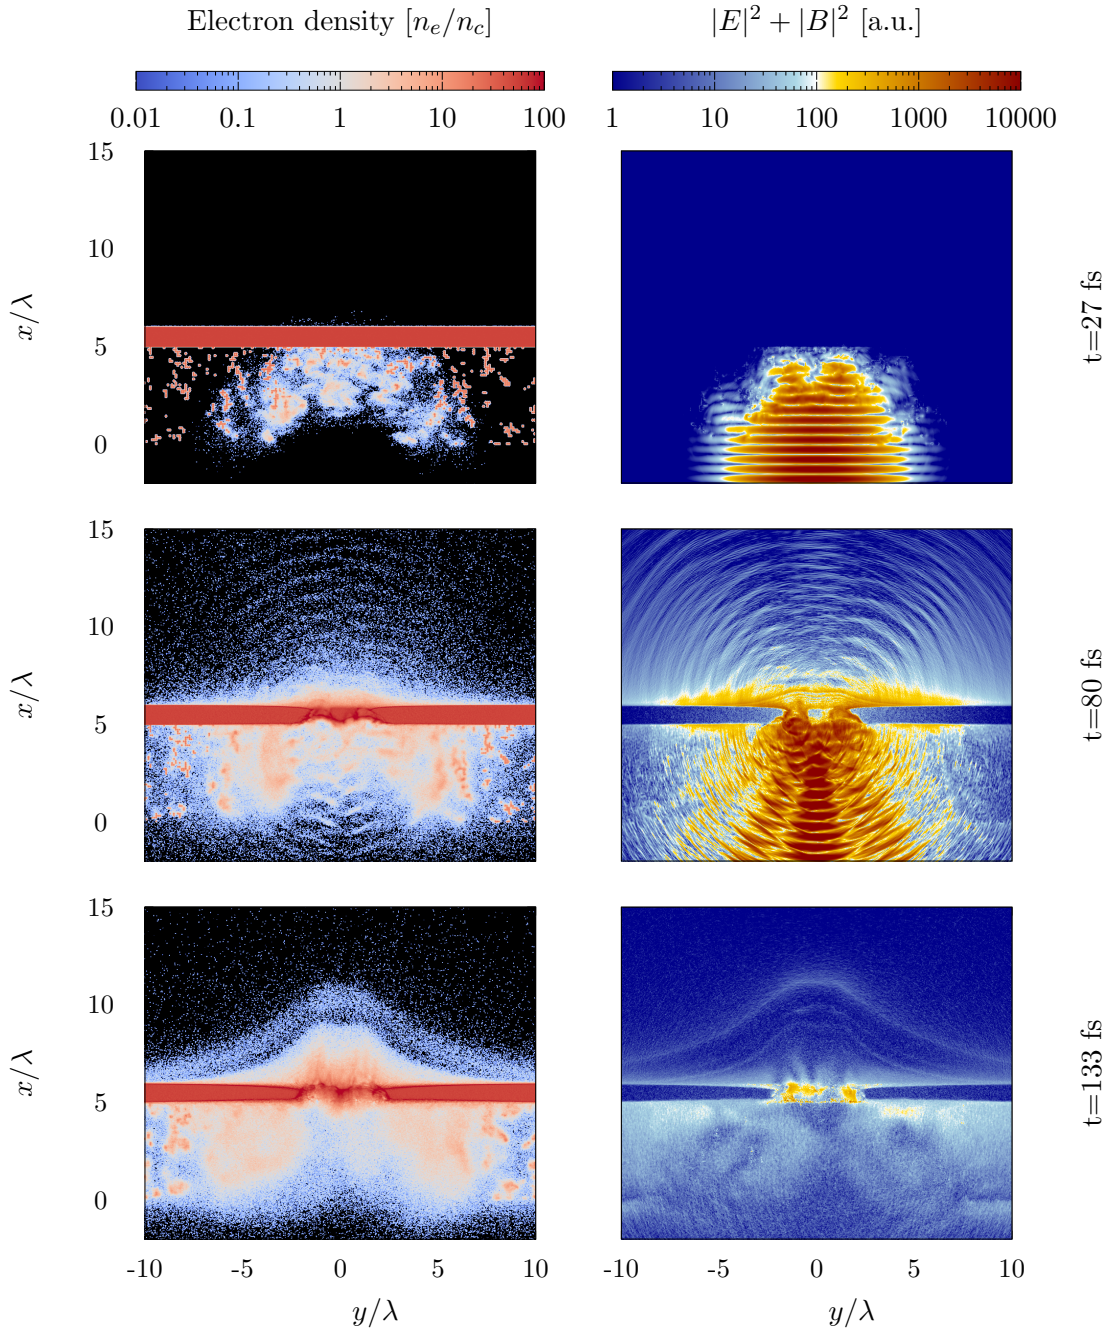


Figure 3.10: Electron density and electromagnetic energy density at plane $z = 0$.

HILL FOAM




Figure 3.11: Electron density and electromagnetic energy density at plane $z = 0$.




Figure 3.12: Time evolution of the energy normalized to the initial laser energy of foam and substrate electrons (left) and contaminant protons (right).

ions follow the electromagnetic field generated by the displaced electrons. Indeed, both foam and substrate electrons are soon heated by the interaction, absorbing a significant fraction of the laser energy (up to 50% with foams) until time $t = 80$ fs (see figure 3.12, left). Part of these electrons acquire enough energy to move across the target and create a cloud located on its rear side, larger when the foam is attached. This strong charge separation generates a high intensity longitudinal electric field that eventually leads to impurity protons acceleration. Such protons gain energy from time $t = 53$ fs and at $t = 80$ fs begins the acceleration process at the expenses of electron energy (see figure 3.12, right). This is indicated also by the evolution of the longitudinal component of the electric field (E_x) that reaches its highest values around 80 fs for the foam-attached targets. In particular, the maximum value of the accelerating field is obtained with the uniform foam (about 50 TV/m), 1.3 times the value obtained with the hill foam. Nevertheless, as mentioned later, the highest proton energies are observed with the hill foam. It is worth mentioning that protons energy has not reached saturation yet (see right panel of figure 3.12), especially with the foam targets. A longer simulation time would likely allow the observation of higher cutoff energies.

The presence of inhomogeneities in the foam strongly changes the characteristics of the interaction. Laser pulse self focusing due to the propagation through a uniform plasma is present with a uniform foam, while strongly reduced with a nanostructured foam because of the density non-uniformities. Moreover, in the hollow foam target the pulse digs two valleys in the solid, while only one in the other cases (see density snapshots 3.8, 3.9, 3.10, 3.11). This may result in a more extended electronic cloud along y direction, thus in a reduced accelerating field. Anyhow, there is no evidence to conclude that this fact is strictly related to the presence of the hollow. It may depend on the particular inhomogeneities of the specific foam. Deeper investigations are needed to clarify this issue.




Figure 3.13: Contaminant protons energy spectra at time $t = 160$ fs, filtered in angles $\theta \in (-5, 5)$, $\phi \in (-5, 5)$

For what concerns the acceleration performances, the foam allows an enhanced energy coupling with the laser pulse with respect to the bare target that results in higher proton energies. This is in agreement with the experimental results in [32] and [33]. The cutoff energies of the contaminant protons are, in ascending order, 12 MeV with no foam, 22.4 MeV with hollow foam, 28.9 MeV with uniform foam and 29.5 MeV with hill foam. The addition of the foam yields to an enhancement of the cutoff energy by a factor between 1.8 and 2.4, not far from those obtained experimentally (see paragraph 1.5.2 and figure 1.15). The proton spectra exhibit the typical exponential-like behaviour also observed in experiments (see figure 1.15) and are very similar for the uniform foam and the foam with the hill (see figure 3.13).

The angular distributions of the contaminant ions with energy around 1 MeV ($\pm 1\%$) are shown in figure 3.15 for the bare and the hill foam-attached targets. Consider the bare solid target. The distribution at time 53 fs is spread in a neighbourhood of the origin, while at time 80 fs (the beginning of the acceleration process) is focused in a small region around the origin, i.e. in a small cylinder around the target normal. Then a ring-shaped distribution with symmetry properties is obtained. Now consider the foam-attached targets (only the hill case is shown, but similar considerations hold for the other foam-attached targets). Until time 133 fs the distribution is almost empty around the origin. At intermediate times it is ring-shaped with non-symmetric spots, then it deforms. At the final time the majority of contaminant ions is located around the origin. The angular distributions of the contaminant ions around higher energies (> 5 MeV) show irregular, filamentous structures. Each case exhibits a different behaviour. In general, high symmetry characteristics are displayed only by the bare solid target and in presence of a foam the distribution is wider. It may happen that there are no protons around the origin (see hill case in figure 3.14). This fact is crucial for experiments because if the ion detector is located along the target normal, many

accelerated ions could be lost in the measurements. As already mentioned in paragraph 1.5, the accelerated ions angular emission profile can be experimentally measured using radiochromic films. Up to now, neither numerical nor experimental RCFs obtained using foam-attached targets can be found in literature. Anyway, both annular and filamentous distributions have been observed in several experiments. Ring-shaped RCFs have been obtained since the early experiments of laser-driven ion acceleration, for example by Clark *et al.* [86]. Lin *et al.* used bare solid targets with different materials, varying the conductivity [87]. They obtained "more filamentous" films with star-like patterns using the lower conductive targets (see figure 3.16). This may suggest that the presence of the foam has somehow the same effects – at least in terms of RCF structures – as a low conductivity bare solid target, for example made of mylar. Filamentous RCFs, i.e. spatially modulated protons, have also been observed by Metzkes *et al.* with μm -thick metal foils [88]. They attribute such structures to electron instabilities that arise depending on laser energy and target thickness.




Figure 3.14: Angular distribution of contaminant protons with energy $E \in [18, 22]$ MeV at the final time $t = 160$ fs. ϕ is the angle between the ion position and the target normal on the (x, y) plane, θ is the angle between the projection of the ion position on the (x, y) plane and z axis.

In conclusion, the acceleration process is enhanced when attaching a foam layer on the front side of a bare solid target. Varying the region hit by the laser pulse, different results have been obtained. It seems that such performances are comparable if the laser pulse hits a region, whether uniform (uniform foam) or inhomogeneous (hill foam), with fixed thickness. Still, in these cases, although the proton cutoff energies and spectra are similar, many differences arise when considering other aspects of the phenomena.

A discussion on the computational resources and burden of the present 3D simulations is presented in section 3.5.

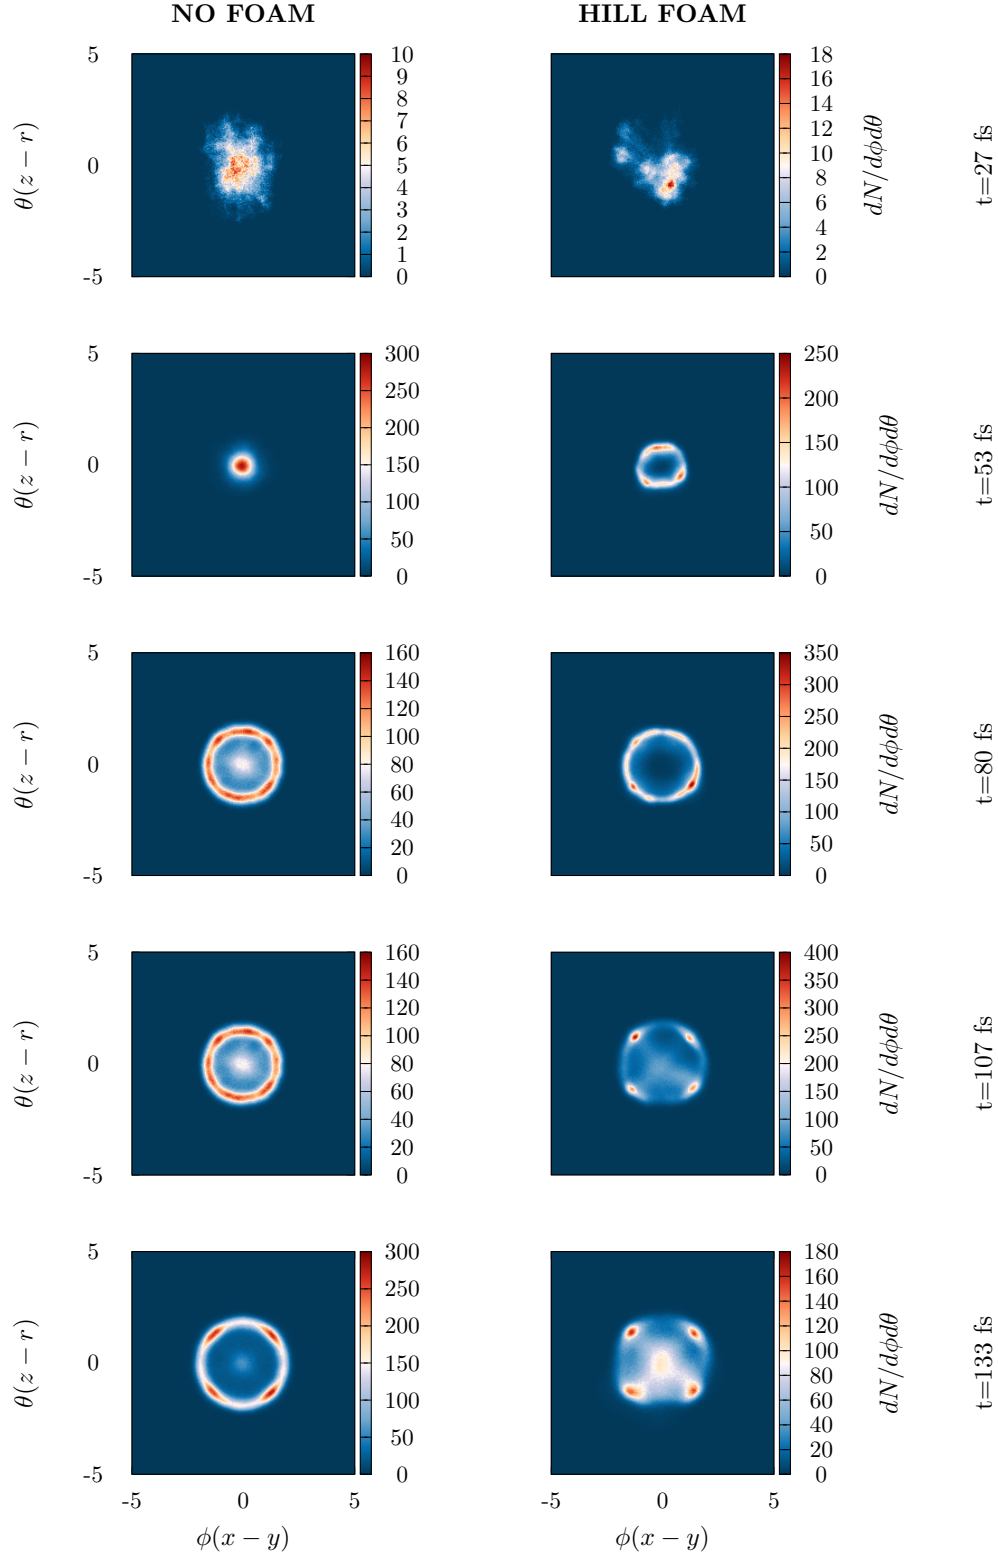


Figure 3.15: Angular distribution of contaminant protons with energy $E \in [0.9, 1.1]$ MeV. ϕ is the angle between the ion position and the target normal on the (x, y) plane, θ is the angle between the projection of the ion position on the (x, y) plane and z axis.

3.4.2 Two-dimensional PIC simulations

A set of two-dimensional simulations has been performed to qualitatively analyse the effect of an idealized nanostructured foam, modelled as a collection of nanospheres regularly arranged in space. A simplified picture is adopted with respect to the previous three-dimensional simulations because of both the reduced dimensionality and the plainer foam model, which allows a wider parametric scan. The simulated targets are:

- bare solid foil,
- uniform foam at critical density attached to a solid foil,
- three different nanostructured foams with average critical density attached to a solid foil.

Each target is irradiated by a high-intensity P polarized Gaussian laser pulse with $a_0 = 10$ at normal incidence. Assuming a wavelength of $0.8 \mu\text{m}$, the peak intensity is $I = 2 \times 10^{20} \text{ W/cm}^2$ ($a_0 = \sqrt{2e^2\lambda^2 I / \pi m_e^2 c^5}$). The laser pulse has a \cos^2 -function temporal profile with full-width-half-maximum of the field equal to 41 fs, equivalent to a full-width-half-maximum of the intensity of 30 fs. The laser focal spot size is $3.2 \mu\text{m}$.

The solid foil is a uniform, high density layer consisting of electrons and ions with $Z/A = 1/2$ (e.g. fully ionized carbon) with $0.4 \mu\text{m}$ thickness. Its electron density is $100n_c$, sampled with 100 electrons per cell, while the ion distribution is sampled with 49 particles per cell. A thin (40 nm-thick) and low density ($10n_c$) contaminant layer is added on the rear side of the solid foil, made of protons and electrons, sampling both populations with 100 particles per cell.

All foams are made of electrons and ions with charge-to-mass ratio equal to $1/2$ (for example C^{6+} ions with $Z=6$, $A=12$) with critical electron density (mean if nanostructured). The uniform foam is a $4 \mu\text{m}$ -thick layer sampled with 4 electrons and 1 ion per cell, respectively. The nanostructured foams are made by dense nanospheres regularly organized in space. Different foams are simulated varying the spheres centre-to-centre distance and radius in such a way that the mean electron density is $1n_c$. The chosen values for these quantities are summarized in table 3.2. The electron density of each sphere is $100n_c$, sampled with 100 electrons, while the ion distribution is sampled with 49 ions per cell. The foam thickness is set to $4\mu\text{m}$.

The simulation time is 187 fs. All species of all plasmas are initialized with a Maxwellian distribution with temperature around 1 eV (10^4 K).

The selection of the simulation parameters (e.g. number of particles per cell) followed the same principles used in the 3D case.

Table 3.3 summarizes the parameters common to all simulations. Figures 3.17, 3.18, 3.19

2D Nanostructured Foams Parameters

	Foam 1	Foam 2	Foam 3
Distance	0.2 μm (0.25 λ)	0.4 μm (0.5 λ)	0.8 μm (1.0 λ)
Radius	11.28nm	22.56nm	45.12nm

Table 3.2: Centre-to-centre distance and radii of the spheres of the nanostructured foams.

and 3.20 display some graphical results of the numerical campaign. The evolution of the electron density (figures 3.17 and 3.18) and magnetic field (figures 3.19 and 3.20) are represented at four time steps for the foam-attached targets.

Many features of the phenomena are similar to those observed with 3D simulations. The laser pulse penetrates into the critical density foam and reaches the solid target, where it is reflected backwards. As the laser propagates into matter, the electron population gains energy and is able to form a cloud on the rear side of the target. Consequently, a strong charge separation occurs (electrons outside the target, ions inside the target) that yields to contaminant protons acceleration.

Although the electron density becomes more uniform as the interaction takes place in all cases, the presence of non uniformities influences the interaction process. Self focusing of the laser pulse only occurs in the propagation through the uniform foam, not with inhomogeneous foams. Non uniformities also affect the electromagnetic field, in particular the magnetic field (B_z is the only relevant component) exhibits interesting structures (figure 3.20).

After about 70 fs, laser energy starts to be converted into electron energy, that in turn is later converted into ion energy (see figure 3.21 and left figure 3.22). A stronger energy absorption by the electron component is observed with foams 1 and 3, but foam 3 leads to a higher conversion of electron energy into ion energy (see figure 3.22, right). However, highest absorption efficiency does not mean highest proton energies. The cutoff energies are 17 MeV, 33 MeV, 34 MeV, 29 MeV, 31 MeV for the no foam, uniform foam, foam 1, foam 2, foam 3 cases respectively (see figure 3.23). As for the 3D case, cutoff energies would probably be higher for longer simulation times, since the protons energy has not reached saturation yet (see figure 3.22, right). Of course the foam with smaller and closer spheres is the most similar to the uniform foam and this is also proved by the similarities in the spectra and cutoff energies. Even though considering highly idealized foam models, this is a positive result, since real foams are made by microscopic clusters with $\sim 10\text{nm}$ radius. However, the results obtained with foam 3 are not far from those obtained with the uniform foam and foam 1. Indeed, in terms of acceleration performances, the least efficient foam




Figure 3.16: RCFs for materials with increasing conductivity from left to right. From [87].

2D Simulation Parameters	
Box size	$160\lambda \times 100\lambda$
Resolution (points per λ)	[300, 287]
Simulation time	70
Solid foil	
Density	$100n_c$
Thickness	0.5λ
Electrons per cell	100
Ions per cell	49
Contaminant layer	
Density	$5n_c$
Thickness	0.05λ
Electrons per cell	100
Protons per cell	100
Laser	
Type	Gaussian
Incidence	0°
Duration FWHM	$15\lambda/c$
Initial peak position	$[-15\lambda, 0]$
Waist	4λ
a_0	10
Foams	
Thickness	5λ
Mean density	$1n_c$

Table 3.3: Summary of two-dimensional simulations setup.


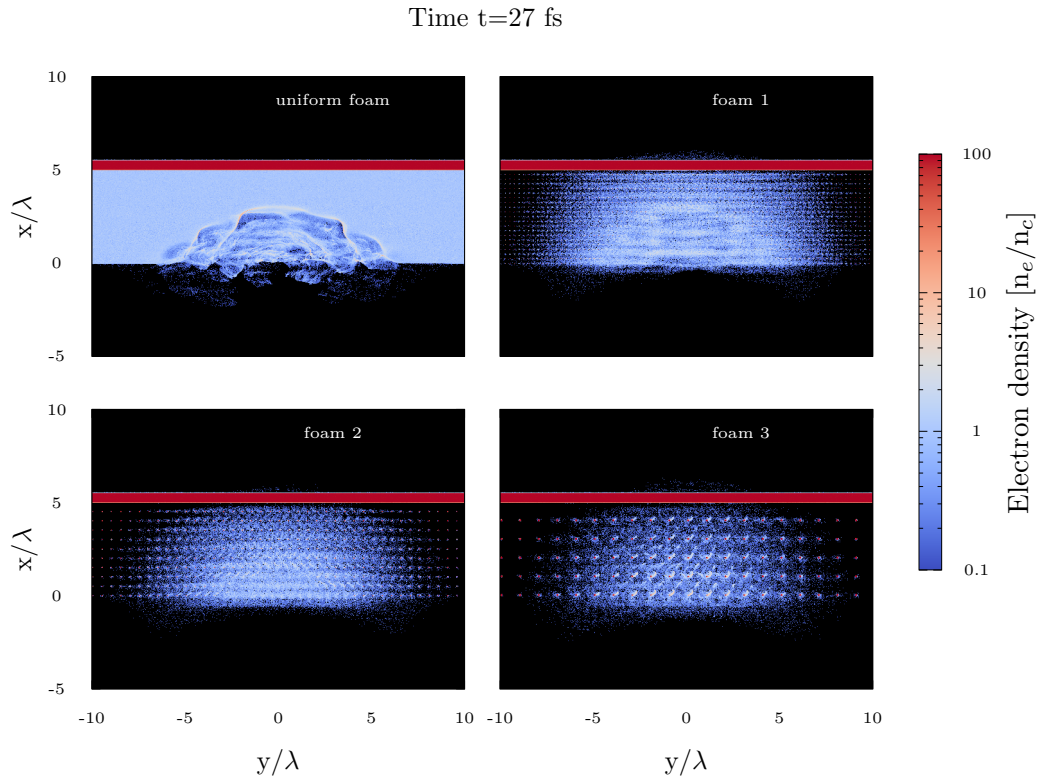


Figure 3.17: Snapshots of the electron density of the foam-attached targets at time 27 fs (top) and 80 fs (bottom).





Figure 3.18: Snapshots of the electron density of the foam-attached targets at time 133 fs (top) and 187 fs (bottom).

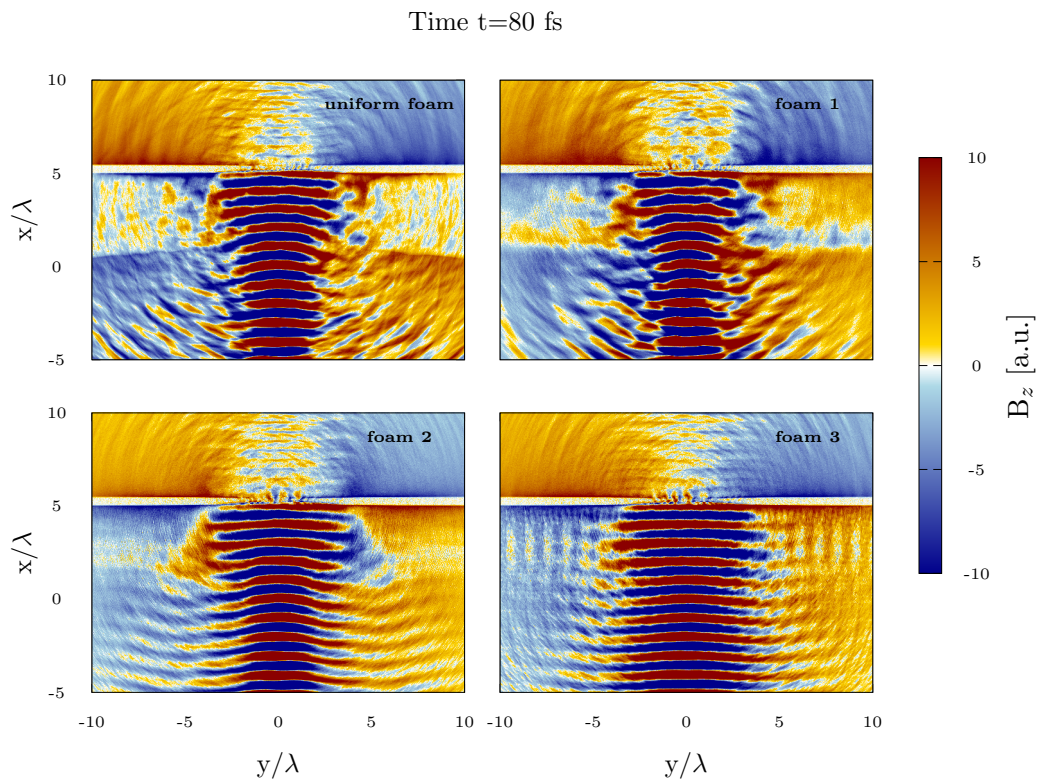



Figure 3.19: Snapshots of the magnetic field B_z at time 27 fs (top) and 80 fs (bottom).

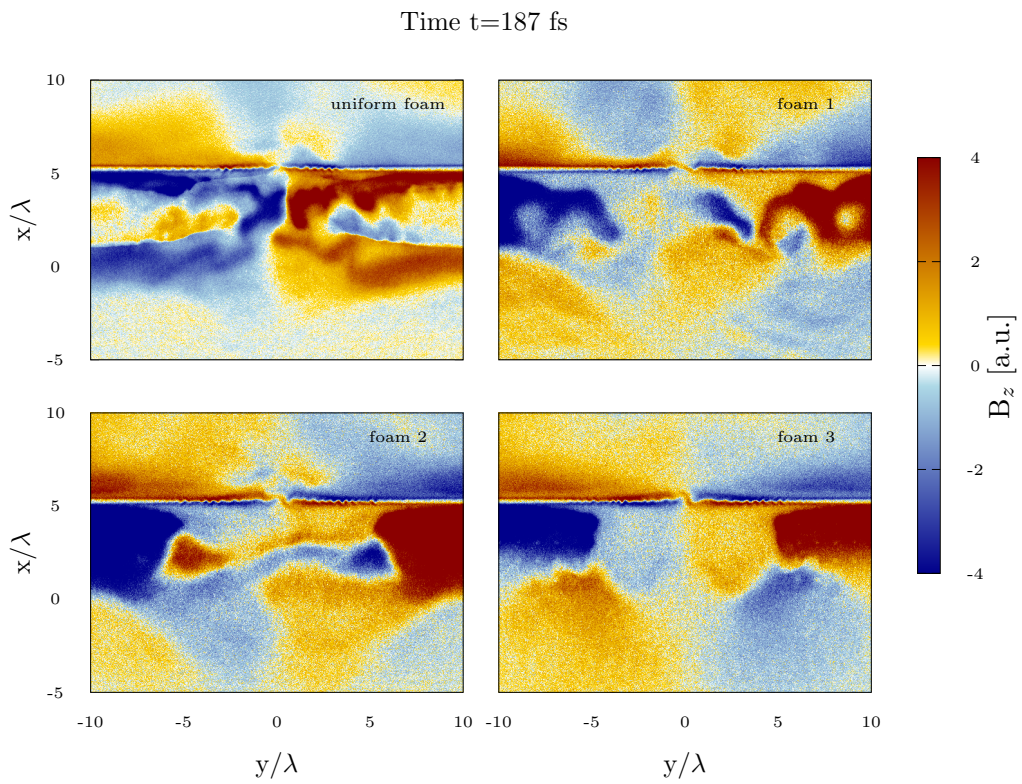



Figure 3.20: Snapshots of the magnetic field B_z at time 133 fs (top) and 187 fs (bottom).

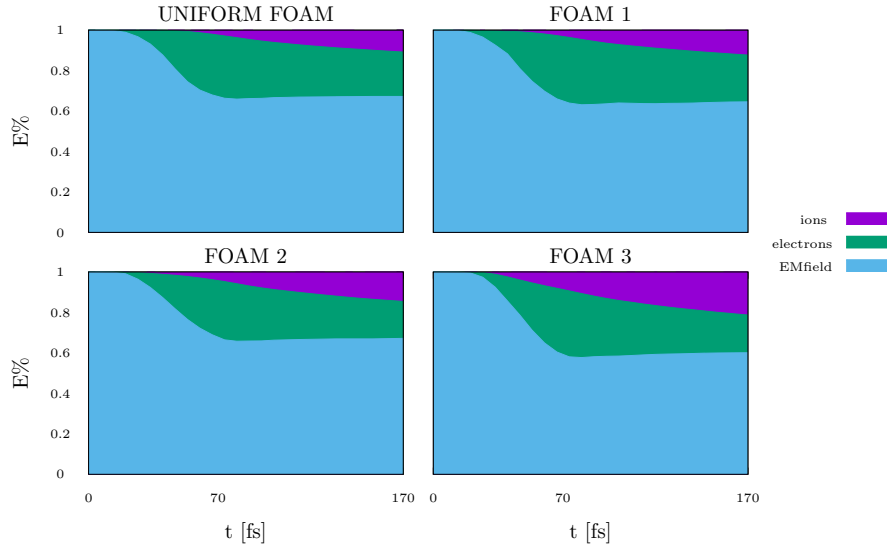


Figure 3.21: Evolution of the energy repartition among electromagnetic field, electrons and ions.

seems to be foam 2, with properties halfway between foam 1 and foam 3. It may be that more favourable conditions are given by the interaction weather with a "quasi-uniform" foam (foam 1) or with distant, larger spheres (foam 3). Further investigations are required to better understand this matter.

Figure 3.24 shows the proton distribution in energy and angle around the time when the acceleration process begins and at the final simulation time. The general shape is similar for all cases, but with smaller and less distant spheres the distribution in angle seems to be wider. A curious spot without particles is obtain at low energies only with foam 1.




Figure 3.22: Time evolution of the energy normalized to the initial laser energy of foam and substrate electrons (left) and contaminant protons (right).




Figure 3.23: Contaminant protons energy spectra at the end of the simulation filtered in angle $\phi \in (-2, 2)$.


3.4.3 Discussion

The foam-attached targets always allow an enhancement of the acceleration process with respect to bare solid targets, irrespective of the simulation dimensionality and whether the foam is uniform or not. This is due to the efficient energy coupling between a near-critical foam and the laser pulse and has already been observed in different experiments. The gain factors for the protons cutoff energies (maximum proton energy with foam/maximum proton energy with simple foil) are between 1.8 and 2.4 for the 3D simulations, between 1.8 and 2 with the 2D simulations. Higher performances are always achieved modelling the foam as a uniform layer, however similar results can be obtained with suitable non-homogeneous foams.

Three-dimensional simulations suggest that the properties of the region hit by the laser pulse play a relevant role. A hollow or a hill in a DLA numerical foam where the laser pulse is expected to irradiate the target leads to different proton energies, higher for the hill case by a factor of 1.3. Furthermore the accelerated protons may be located on a wide solid angle around the target normal. Hence the "thickness" of the illuminated spot strongly affects the acceleration performances. Comparisons with experimental results suggest that these kinds of simulations may be predictive.

Two-dimensional simulations suggest that the dimension and distance of the clusters composing the foam can lead to different results. The highest cutoff energies are obtained with the uniform foam and with the smallest and closest nanospheres. With intermediate values of the spheres radius and centre-to-centre distance the acceleration process seems to be less efficient.

Proton distribution in energy and angle at time $t=80$ fs



Proton distribution in energy and angle at time $t=187$ fs

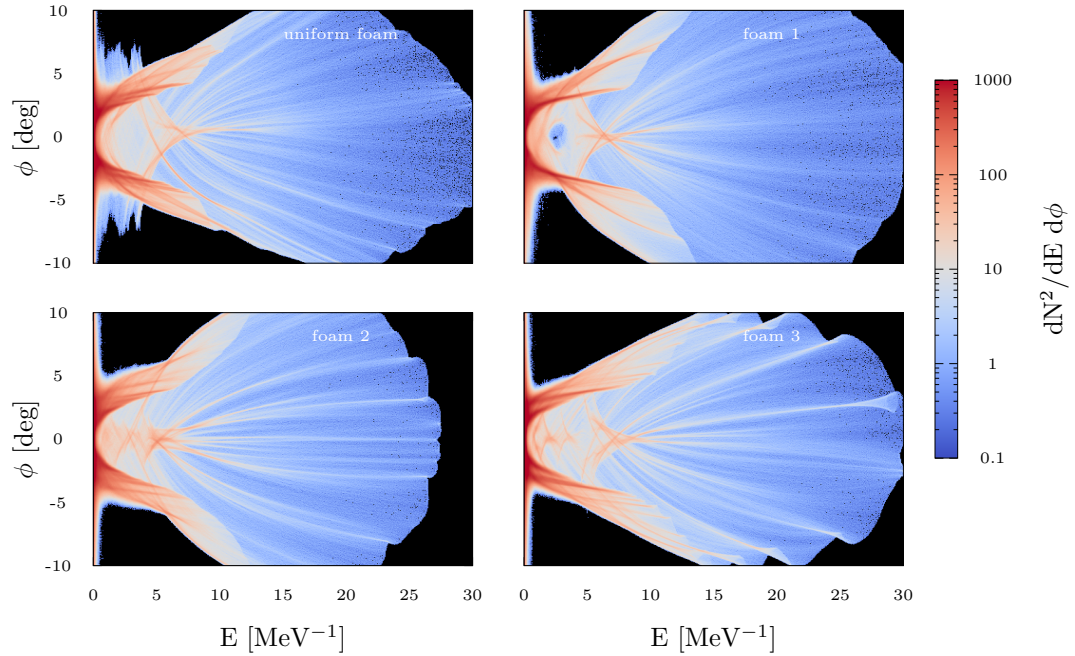


Figure 3.24: Distribution in angle and energy of the contaminant protons at two time steps.

Therefore, the acceleration performances are influenced by many properties of the foam, such as non-uniformities characteristics, thickness, clusters size and cluster-to-cluster distance.

3.5 Computational issues

The two-dimensional simulations presented in the previous section required up to 40×10^3 CPUhours on Fermi supercomputer each, while the three-dimensional simulations required a much higher computational cost, up to 70×10^3 CPUhours on Marconi supercomputer each (1 Marconi CPUh \approx 7 Fermi CPUh). Consider that, typically, a budget of up to few million Fermi core hours can be awarded by a Cineca call for high performance resources. Under such conditions, only few 3D simulations like those described in the precedent section could be afforded.

The computational cost of a three-dimensional numerical simulation can be extremely high, both in terms of memory occupation and CPU hours. Moreover, the only feasible 3D simulations require sets of parameters rather far from those actually used in the experiments of laser-plasma interaction. One must accept several compromises when setting a 3D (and sometimes also 2D) simulation, solely for economy reasons, what is more bearing in mind that a high enough resolution is needed. Indeed, in numerical experiments of laser-plasma interaction it is crucial to solve the skin depth (see section 1.4), i.e. one of the main characteristic scale length. Because of this together with the Courant condition, only short simulation times are possible, about hundreds of femtoseconds. Furthermore real solid-density plasmas have electron densities of $200n_c$, while the corresponding numerical values are about $40 - 80n_c$. This is due to the fact that sampling a higher electron density requires a higher number of particles per cell, hence the choice falls on the minimum density value that allows the simulation of an overdense plasma (i.e. a plasma that reflects the laser pulse). Besides, allowed foam parameters are limited. For instance the radius of numerical nanospheres is 40-50 nm, while real nanoparticles have radii around tens of nm. Moreover, real foams can have thickness up to tens of micrometers, while at most $10 \div 30\mu\text{m}$ -thick foams can be simulated in 3D.

In general, a high number of particles per cell (PPC) is required to limit the numerical noise due to the sampling of phase space (numerical noise scales as $\text{PPC}^{-1/2}$). If n_e is the plasma electron density, as a rule of thumb at least n_e/n_c PPC (and no less than 1 PPC) are required. Another important issue is related to the resolution. To be more specific, given the number of cores and the physical scenario (box size, simulation time, plasma species...), the computational time needed to carry out the simulation with a standard

PIC code scales as

$$T_{comp} \sim \alpha \left(\frac{1}{\Delta}\right)^{n+1} + \beta \text{PPC} \left(\frac{1}{\Delta}\right)^{n+1}$$

where Δ is the grid cell size, n is the dimensionality of the simulation and α, β suitable constants ("+" at the exponent is a consequence of the Courant condition). The two terms are related to the computations concerning the electromagnetic field and the particles, respectively. Usually $\text{PPC} \gg 1$, hence the particles computations give the dominant contribution to the computational time. Given the resolution Δ and the number of particles per cell PPC , a 3D simulation would require at least Δ^{-1} times the time required for the corresponding 2D simulation, that is a lot more. Not to mention that this is a strong underestimation, since for example 2D current deposition involves 9 grid points, while 27 in 3D. Affordable resolutions in three-dimensional simulations do not allow to solve the nanoscales. This means, for example, that the smallest inhomogeneities scales of the foam targets are completely neglected. This may not be a crucial fact if the laser wavelength is much greater than such scale lengths, but still represents a constraint in the choice of the parameters.

Extensive numerical campaigns able to give valuable quantitative information (i.e. 3D) require extremely high computational cost and are often unaffordable. Reduced 2D simulations are way more feasible than 3D simulations, but are not able to capture some important features such as the real morphology and structure of the foam layers or the inherently 3D physical effects. Therefore, optimized code implementation is crucial to perform, above all 3D, simulations of real systems. Chapter 4 deals with some aspects related to this issue.

Chapter 4

Code improvements

This chapter presents some extensions and improvements developed in the code *piccante*. In particular, section 4.1 presents the strategies adopted to add OpenMP parallelism into *piccante* along with their testing, making the code hybrid. Following, section 4.2 deals with the implementation and testing of a charge-conserving current deposition algorithm, i.e. the Esirkepov algorithm, alternative to the standard method that straightly computes the current from the shape functions. A summary of these activities follows in section 4.3.

4.1 OpenMP parallelism

As already discussed in section 3.5, reliable simulations of systems of actual interest for research purposes require a great computational cost, namely a high space-time resolution together with a high number of macro-particles (often $10^7 \div 10^9$), hence they can only be performed on computer clusters. This implies the exploitation of parallel programming models, among which MPI [89] and OpenMP [90] are the most widely used. MPI defines a standard for message passing programs. It is designed for distributed memory systems (e.g. the cores of a CPU) where a program is executed by a set of *tasks* that do not share the same memory regions, but can communicate with each other. The programmer has to explicitly implement all parallelisms. On the other hand, OpenMP is designed for sharing memory systems where a program is executed by a set of *threads* that globally share the same memory. It is easily implemented by a set of compiler directives, library routines and environment variables. Large supercomputers always require the exploitation of at least the MPI model, since the code runs on multiple CPUs. For what concerns particle-in-cell codes, efficient hybrid parallelization is still an open research topic, continuously evolving as a consequence of the non-stopping development of novel architectures (see [91], [92], [93]).

The code *piccante* already comes with MPI parallelization. Here it is reported on the development of the OpenMP model in *piccante*, which makes the code hybrid.

It is well known, and can be immediately verified with a profiling analysis, that a PIC code spends most of the execution time ($\gtrsim 90\%$) on computations related to the particles, namely projecting the current density field on the spatial grid and updating the particles momenta. The OpenMP parallelization focused on these portions of the code. Furthermore, to avoid excessive overhead, the parallelization also concerned the routine responsible for the enforcement of the periodic boundary conditions on the particles.

4.1.1 Parallelization strategies

Here the main OpenMP optimization strategies developed into the code *piccante* are discussed. The pseudocodes of the original routines involved in the parallelization are presented, followed by a discussion on the adopted strategy together with some code excerpts. It is worth pointing out that the basic data structure used in *piccante* consists essentially in arrays. For instance particles data (positions, momenta,...) and fields (electric, magnetic, current density) are all stored in arrays of proper size.

Current deposition

To evaluate code performances the standard current deposition method has been chosen over the Esirkepov method because it requires less computational resources. The pseudocode for the standard current deposition routine is the following.

```

for each particle do
  for each spatial dimension do
    half-advance particle position;
    find nearest integer and half-integer nodes and compute corresponding
    positions on the grid;
    compute weights on nearest integer and half-integer nodes and two adjacent
    nodes;
  end
  for each of the nodes involved in the deposition do
    compute particle contribution to current density;
    add the particle contribution to total current density;
  end
end

```

It is important to note that every particle is independent from the others, thus the computation of quantities related to different particles can happen simultaneously. Nevertheless, the computation of the total current should be serial, since several particles can deposit their contribution on the same nodes, i.e. write in the same elements of the total current array. Two different parallelization strategies have been developed, labelled as "use of atomics" and "auxiliary currents" respectively.

Use of atomics The first, simplest, approach consists in adding some `#pragma` directives to the code. In particular, a `#pragma omp parallel for` directive is added before the loop on the particles. Such directive specifies that the iterations of the following loop must be executed simultaneously by a team of threads, which are created in that moment. The strategy used to carry out the parallel loop can be partially controlled using some specific clauses (e.g. `private`, `shared`, `reduction`, `schedule`,...). In particular, the `schedule` clause describes how iterations are divided among the different threads. Three available options are: `static` if every thread carries out the same amount of iterations; `dynamic` if when a thread finishes one chunk, it is dynamically assigned to another; `auto` if the decision is entitled to the compiler and/or runtime system. Static scheduling is not recommended because evenly divided loop iterations may not require the same amount of time. Therefore, it is better to opt for a dynamic scheduling. Even better is to let the scheduling type be automatically determined because the optimal chunk size is not known a priori. This is why, after some tests, the clause `schedule(auto)` has been adopted. All the variables related to the single particles need to be specified in a `private` clause. Three `#pragma omp atomic` statements are added before the portions of the code where each component of the total current array is computed, so that its memory location is updated by only one thread at a time. In fact, the `atomic` directive is used to serialize a specific operation so that only one thread is allowed to update the related memory location. Atomic operations usually use hardware instructions, hence come with relatively low overhead, but the only presence of such operations can significantly affect the optimization. In the following, some code excerpts are reported.

```
#pragma omp parallel for schedule(auto) private(p, gamma_i, i, j, i1, j1, k1, i2, 1
    j2, k2, hii, wii, hiw, wii, rr, rh, rr2, rh2, dvol, xx, vv)
for(p = 1; p < Np; p++){ 2
// computation of the single particle contribution JXP, JYP, JZP to the 3
    total current JX, JY, JZ, where JXP, JYP, JZP are the right-hand-sides
    of the following lines
#pragma omp atomic 4
*jX += weight*hiw[0][i]*wiw[1][j]*vv[0]*chargeSign; 5
#pragma omp atomic 6
*jY += weight*wiw[0][i]*hiw[1][j]*vv[1]*chargeSign; 7
#pragma omp atomic 8
*jZ += weight*wiw[0][i]*wiw[1][j]*vv[2]*chargeSign; 9
```

}

10

In order to further optimize the performances, data locality is beneficial because it helps to reduce cache misses. This can be achieved trying to make each thread responsible for a set of particles that need to deposit their current in regions which are close in memory, i.e. nearby on the grid. For this reason the particle vector is rearranged, performing a partial sorting based on the particles' x -coordinates. The simulation box is divided along the x axis in as many regions as the number of threads. The particles with x -coordinates belonging to the same region are moved so to be located contiguously in the particle vector, but not in order (it is not a proper sorting). It is a trade-off between data locality and parallelism, since particle partition and reordering must be executed serially before depositing the current in the main loop. These operations, which are carried out serially, are performed by two functions that are declared as follows.

```
int SPECIE::particle_partition(double pivotValue, int left, int right); 1
void SPECIE::particle_reordering(GRID * grid); 2
```

The first of these routines performs the partition of the particle vector. It groups the vector (from index left to index right) into two parts, those less than a pivot value, and those greater or equal to the pivot value. The second routine groups together those particles with x -coordinate belonging to the same region (identified by a specific thread). These operations always resulted convenient if performed right after the particles initialization, i.e. at the first time step. It could be advantageous to repeat the reordering with an a priori chosen frequency, that is inevitably problem-dependent. Anyway, if the particles remain approximately in the same region for all the evolution it may be enough to perform it only once, at the beginning of the temporal cycle. In general, there is a trade-off between the additional cost of reordering operations e related benefits.

Auxiliary currents The second strategy consists in defining as many global copies of the total current array as the number of threads at the first time step. These copies are initialized to zero, filled separately by each thread and then added together in a parallel loop. In this way it is still possible to parallelize the loop over the particles in the current deposition routine with a `#pragma omp parallel for`, with the advantage that no atomic operations are needed. The obvious drawback is the increase of the memory occupation at most by a factor of $3N_{threads}/(9 + 7PPC)$, which is often affordable. This overestimate can be deduced for a simulation box completely filled with particles considering that for every time step 9 double numbers for each grid point are stored ($E_x, E_y, E_z, B_x, B_y, B_z, J_x, J_y, J_z$) together with 7 double numbers per particle ($x, y, z, p_x, p_y, p_z, \text{weight}$). For example, if the simulation box is completely filled with particles, for 16 threads and 100 PPC it is a 7% increment of the memory usage.

The auxiliary currents are added in the class `CURRENT` together with the corresponding accessors methods. They are public since they need to be modified by other methods of the class `SPECIE`.

```

class CURRENT {
public:
double **JXaux, **JYaux, **JZaux;
double **getJXaux, **getJYaux(), **getJZaux();
// ...
}

```

The loop for the deposition looks like as follows.

```

#pragma omp parallel for schedule(auto) private(p, gamma_i, i, j, i1, j1, k1, i2,
        j2, k2, hii, wii, hiw, wiw, rr, rh, rr2, rh2, dvol, xx, vv)
for(p = 1; p < Np; p++){
const int IDthread = omp_get_thread_num();
// for every node involved in the deposition: computation of the single
particle contribution JXP, JYP, JZP to the total current JX, JY, JZ
where JXP, JYP, JZP are the right-hand-sides of the following lines
myJXaux[IDthread][my_indice(edge, 1, 0, 0, i2, j1, k1, N_grid[0], N_grid[1], N_grid
    [2], 1)] += weight*hiw[0][i]*wiw[1][j]*vv[0]*chargeSign;
myJYaux[IDthread][my_indice(edge, 1, 0, 0, i1, j2, k1, N_grid[0], N_grid[1], N_grid
    [2], 1)] += weight*wiw[0][i]*hiw[1][j]*vv[1]*chargeSign;
myJZaux[IDthread][my_indice(edge, 1, 0, 0, i1, j1, k2, N_grid[0], N_grid[1], N_grid
    [2], 1)] += weight*wiw[0][i]*wiw[1][j]*vv[2]*chargeSign;
}
}

```

Momenta advance

The Boris pusher algorithm explained in paragraph 2.2.3 is used to update the particles momenta. The pseudocode is the following.

```

for each particle do
    for each spatial dimension do
        find nearest integer and half-integer nodes and compute corresponding
        positions on the grid;
        compute weights of the particle on nearest integer and half-integer nodes and
        two adjacent nodes;
        compute force acting on current particle interpolating the field values;
        update particle momenta;
    end
end

```

Parallelizing this routine is straightforward because there are no dependencies inside the loop, i.e. every iteration is independent from the others. It is enough to force the iterations of the loop over the particles to be executed in parallel by a team of threads by adding a `#pragma omp parallel for` directive. Again an automatic scheduling policy is adopted for the same reasons as in the case of the current deposition routine. All the variables related to the particles are declared to be private to each thread. The code looks like as follows.

```
#pragma omp parallel for schedule(auto) private(p,c,gamma_i,hii,wii,hiw,wiw 1
    ,rr,rh,rr2,rh2,dvol,xx,E,B,u_plus,u_minus,u_prime,tee,ess,dummy,i,i1,j1
    ,k1,i2,j2,k2)
for (p = 0; p < Np; p++){ 2
// computations... 3
} 4
```

Periodic boundary conditions for particles

After the parallelization of the current deposition and momenta advance routines was completed, different profiling analysis revealed the enforcement of the boundary conditions on the particles to be the most important bottleneck when considering a pure OpenMP context; on the contrary this routine is very efficient for pure MPI parallelization. Optimization of the related source function was then necessary to be able to compare pure OpenMP with pure MPI schemes. Indeed, in this case it is important to take into account the already existing MPI parallelization. In fact, in a multi-task execution the spatial grid is sliced in regions, each one assigned to a different task. When a particle moves from one of these regions to another, the corresponding tasks need to communicate to each other particle's data. The routine that deals with these issues also takes care of the imposition of the particles periodic boundary conditions. The pseudocode for every task prior to OpenMP parallelization is the following. The loop over the particles has loop-carried dependencies, hence a proper parallelization using the `#pragma omp parallel for` directive is possible as long as some arrangements are made. To be more specific, the buffers of outgoing particles have to be private variables and merged together after the cycle. Moreover, the number of total lost particles (`nlost`) must be incremented by only one thread at a time in a critical section. Lastly the shifting of the particles array elements cannot remain inside the loop. Hence, the number of lost particles both along right and left directions together with two buffers containing the corresponding particles' data are defined for each thread as follows.

```
int * nright_loc, * nleft_loc; 1
nright_loc = new int [Nthreads]; 2
nleft_loc = new int [Nthreads]; 3
4
```

```

for each spatial dimension do
    for each particle do
        if particle exits domain across right or left border then
            increment total number of lost particles;
            if particle exits total domain then
                enforce periodic boundary conditions;
            end
            save particle in right or left buffer;
        else
            shift backwards current particle in the particles vector by the amount of
            lost particles;
        end
    end
    send buffers of lost particles to appropriate tasks;
    receive buffers of incoming particles;
    update number of particles;
end

```

```

double **sendr_buffer_loc, **sendl_buffer_loc;           5
sendr_buffer_loc = (double**) malloc(Nthreads*sizeof(double*)); 6
sendl_buffer_loc = (double**) malloc(Nthreads*sizeof(double*)); 7

```

Moreover, a vector containing all the indexes of lost particles is defined as:

```

std::vector<int> pIndexLost;                               1

```

and is used to rearrange the particles vector once the loop has been completed, as explained later. These variables are declared as shared among the threads and the loop is parallelized as follows.

```

#pragma omp parallel for schedule(auto) shared(nlost,nright_loc,nleft_loc, 1
    sendr_buffer_loc,sendl_buffer_loc,pIndexLost)
for (p = 0; p < Np; p++){                                  2
    const int IDthread = omp_get_thread_num();              3
    // computations...                                     4
}                                                            5

```

The variable `IDthread` is the unique thread number used to handle the current iteration of the parallel loop. As soon as a particle is found to be exiting the current sub-domain, the total number of lost particles is incremented by one and the particle index is saved in `pIndexLost`, all in a critical section. If the particle exits across the right (left) border, `nright_loc[IDthread]` (`nleft_loc[IDthread]`) is incremented by one and its data is added at the end of the `sendr_buffer_loc[IDthread]` (`sendl_buffer_loc[IDthread]`)

vector. Once the loop has been carried out, the vector containing the data of the particles assigned to the current task has to be rearranged in such a way that the remaining particles are saved in its front (the leaving particles can be overwritten since they are saved in the buffers). To efficiently do so, the remaining particles saved in the back of the particles vector are copied in the elements corresponding to the leaving particles saved in the front of the same vector (see figure 4.1). Now the current task is ready to send the buffers of leaving particles and receiving the buffers of incoming particles. The numbers of incoming and leaving particles are used to compute the updated number of total particles (equal to: old total number of particles - number of leaving particles + number of incoming particles). At this point the particles vector can be reallocated, being lengthened or shortened whether there are more incoming particles than leaving or not; in any case, incoming particles data are added to its back.

It is worth mentioning that in many situations, especially in those employed for testing code performances, the number of particles leaving each sub-domain is small (few particles per task). The significant overhead due to this routine observed during profiling suggests that the most burdensome operations consist in the conditional statements that check the coordinates of every single particles to establish whether they are leaving or remaining. Therefore the adding of the `#pragma omp critical` sections does not cause too much overhead, being executed only when a leaving particle is found.




Figure 4.1: Scheme of the algorithm used to rearrange the array containing particles' data. All remaining particles need to be stored in its front. Starting from the back, as soon as a remaining particle is found, its data is copied in the position corresponding to the first leaving particle, and so on.

4.1.2 Numerical tests

The code has been tested on Galileo IBM NeXtScale cluster hosted at CINECA, equipped with 2 octa-core Intel(R) Xeon(R) CPU E5-2630 v3 @ 2.40GHz per compute node. The Intel compiler has been used with the following optimizations: `-O3 -std=c++11 -xCORE`

`-AVX2 -ipo -openmp`. Extensive profiling work has been carried out with Intel VTune Amplifier 2016 software. Two different configurations have been probed both with a uniform, low-temperature plasma, no laser sources. In the first case the simulation box is only partially filled with plasma, while it is completely filled in the second case. Two different reasons led to the choice of such "simple" configurations. First, the will to avoid numerical and physical instabilities to occur. Second, the need to simulate situations dominated by the particles dynamics, since the most onerous routines are precisely the ones involving particles computations. A completely serial execution of the code shows in both cases that about 40 – 60% of the computational time is spent in the current deposition routine and about 20 – 40% in the momenta advance routine. Hence, these functions have been parallelized first, according to the strategies explained in section 4.1.1. An initial set of tests has been performed turning off the non-particle-related portions of the code. These tests showed a not negligible overhead caused by the routine responsible for the enforcement of the boundary conditions on the particles positions. After parallelizing this third routine, satisfactory performances could be achieved, even though they – unsurprisingly – proved to be problem-dependent. The simulations setup is summarized in table 4.1; to easily visualize the two configurations see figure 4.2.

Partially filled box Plasma size and position have been chosen in order to compare the performances of pure OpenMP parallelization with pure MPI parallelization in a very particular scenario, working with 16 threads vs. 16 tasks. Using only MPI the box is sliced in 4×4 equal regions, each one assigned to a different task. This means that only one task is responsible for the sub-domain occupied by the plasma, therefore the load is highly unbalanced. The other 15 tasks are assigned to regions where the plasma is absent, so they need to perform a lot less computations, i.e. only field-related computations. On the other hand, when using OpenMP, a single thread is not bounded to a particular region of the spatial grid. On the contrary, all threads of the team work simultaneously on the whole grid. Under these conditions the OpenMP model is expected to show the best performances. Of course this is an ad hoc situation, but something similar could occur in simulations of physical interest with plasmas characterized by particular strong non-uniformities, as explained later. Table 4.2 summarizes the main results.

OpenMP seems to work much better than MPI with the localized plasma, which was expected. The first method is more than 4 times faster than MPI. With both strategies, profiling analysis showed that for the most time all 16 threads are working efficiently with no potential gain in the parallel regions and that the load balance is ideal. Between the two deposition strategies, the "use of atomics" showed better performances. This means that the cost of the atomic operations is low if compared to the cost of allocating, computing and adding together the auxiliary copies of the current density vector. This is true because the number of threads is quite big, which makes the "auxiliary currents" strategy less

Simulation Parameters	
Box	$[-50, 50] \times [-50, 50]$
Grid points	2000×2000
Time steps	298
Courant factor	0.98
Field boundary conditions	periodic
Particles boundary conditions	periodic
Electrons per cell	5×5
Ions per cell	5×5
Partially filled box	
Plasma	$[0, 25] \times [0, 25]$
Total number of particles	12.5 millions
Completely filled box	
Plasma	$[-50, 50] \times [-50, 50]$
Total number of particles	0.2 billions

Table 4.1: Summary of the simulations setup for the evaluation of code performances with OpenMP parallelization.




Figure 4.2: Partially filled box vs. completely filled box.

Code	Tasks	Threads	Time [s]
Use of atomics	1	16	161.496
Auxiliary currents	1	16	559.012
MPI	16	1	774.251

Table 4.2: Performance results for partially filled box test.

efficient. Left panel of figure 4.3 shows the scaling of the computational time with the number of threads. While using the atomic operations the scaling is satisfactory, with the auxiliary currents it is not. This is not surprising because in this case, as the number of threads increases, also the number of copies of the current vector increases. A possible improvement could consist in defining the size of the current vector copies in relation to the number of threads, so to save time and memory. Every thread would work on a portion of the grid, not fixed a priori. This feature will be implemented in the future.

Completely filled box In this scenario the best performances are expected from the MPI model, since it is known to work better than OpenMP when data size is moderate and the problem is computation-intensive. The aim here was to get OpenMP performances as similar as possible to those of MPI. The best that could be accomplish was a computational time with OpenMP $1.6\times$ than with MPI, working again with 16 threads vs. 16 tasks. Table 4.3 summarizes the main results. The situation is very different from the previous

Code	Tasks	Threads	Time [s]
Use of atomics	1	16	1979.659
Auxiliary currents	1	16	1386.623
MPI	16	1	857.050

Table 4.3: Performance results for completely filled box test.

case. The "auxiliary currents" strategy, even if not optimal in its present implementation, performs better than the "use of atomics" strategy. This is because the atomic operation of the latter cause excessive overhead, as opposed to the case of the partially filled box. The scaling of the computational time with the number of threads is acceptable (see right panel of figure 4.3), but MPI is always faster.

As a remark, these results can be exploited to estimate the execution time per particle per time step as

$$\frac{\textit{elapsed time} \times \textit{number of tasks/threads}}{\textit{number of time steps} \times \textit{number of particles}}$$

which results equal to 230 nanoseconds for MPI and 372 nanoseconds for OpenMP.




Figure 4.3: Partially filled box vs. completely filled box.

Other tests As already mentioned, the OpenMP parallelization could be convenient when simulating plasmas with special spatial inhomogeneities. For instance, if the plasma is highly localized along one particular direction, suppose z , then one could use MPI on the (x, y) plane and OpenMP along z (see figure 4.4). Based on this idea, the two-dimensional configuration depicted in figure 4.5 has been tested, where a plane wave is irradiated onto a stick-like plasma. A first simulation was performed using a pure MPI parallelization, slicing the box into 4 and 36 sub-domains along the x and y directions respectively. The obtained elapsed time is equal to 490 seconds. A second simulation was performed using an hybrid scheme, slicing the box in 36 regions along y direction each related to a single MPI task and generating 4 OpenMP threads using the "use of atomics" strategy. In this case the elapsed time is 325 seconds, with a gain factor of 1.5. It is worth mentioning that great interest has been recently addressed towards these kinds of plasma, such as in [94].

Lastly, OpenMP parallelization can be extremely useful when running the code on multi-core architectures, such as Intel Many Integrated Core Architecture machines (MIC). A preliminary set of tests has been performed specifically on Xeon Phi Coprocessor 7120P (16GB, 1.238 GHz, 61 core), using a hybrid model. A Xeon Phi coprocessor has 60 cores on shared memory with Linux on board. It has simplified architecture, if compared to a CPU core, and lower peak performances. Nevertheless, larger numbers of Xeon Phi cores can reach about few TFLOPS peak performances in double precision. They are equipped with 512-bit vector registers and support up to 4 threads per core. Using the same simulation setups as in the partially and completely filled box cases, optimum performances seem to be obtained with 3 MPI tasks and 80 OpenMP threads, using the corresponding optimal strategy observed on regular processors. Further tests to confirm these promising results




Figure 4.4: Snapshot of a simulation with a strongly non-uniform plasma, localized mainly along one direction (red axis). OpenMP parallelization along that direction together with MPI parallelization along the other two coordinates (yellow and green axes) may result convenient.




Figure 4.5: Initial condition with strongly non-uniform plasma for OpenMP performance testing. An hybrid execution showed a 1.5 gain in the elapsed time with respect to the purely MPI case.

are needed, especially because the code will be exploited in several numerical campaigns on the CINECA supercomputer Marconi, that will be soon equipped with the latest generation of MIC architectures.

4.2 Charge-conserving current deposition method

A pivotal part of every PIC code is the current deposition method. Depending on the problem under consideration, different features may be desired. Efficiency may be priority in large scale simulations, while accuracy may be preferred when investigating specific details of the physics. In general, a trade-off between the two is unavoidable.

The standard algorithm for current deposition is based on the projection of the particles positions and velocities onto the grid via shape functions. The scheme for computing \mathbf{J} is condensed in formula (2.4). When using a second order FDTD Maxwell solver on a Yee-lattice the same formula in expanded form reads (dropping species index):

$$\begin{aligned} J_{x_{i+\frac{1}{2},j,k}}^{n+1/2} &= \frac{q}{\Delta x \Delta y \Delta z} \sum_{p=1}^{N_p} v_{xp} B_{i+\frac{1}{2}}(x_p^{n+1/2}) B_j(y_p^{n+1/2}) B_k(z_p^{n+1/2}) \\ J_{y_{i,j+\frac{1}{2},k}}^{n+1/2} &= \frac{q}{\Delta x \Delta y \Delta z} \sum_{p=1}^{N_p} v_{yp} B_i(x_p^{n+1/2}) B_{j+\frac{1}{2}}(y_p^{n+1/2}) B_k(z_p^{n+1/2}) \\ J_{z_{i,j,k+\frac{1}{2}}}^{n+1/2} &= \frac{q}{\Delta x \Delta y \Delta z} \sum_{p=1}^{N_p} v_{zp} B_i(x_p^{n+1/2}) B_j(y_p^{n+1/2}) B_{k+\frac{1}{2}}(z_p^{n+1/2}), \end{aligned}$$

The analogous approximation used for the charge density ρ is the following:

$$\rho_{i,j,k}^n = \frac{q}{\Delta x \Delta y \Delta z} \sum_{p=1}^{N_p} B_i(x_p^n) B_j(y_p^n) B_k(z_p^n), \quad (4.1)$$

where ρ and \mathbf{J} are computed on grid nodes and time steps consistently with the FDTD Maxwell solver. The main drawback of this approach is the resulting non-conservation of the total charge, meaning that the continuity equation, in general, is not satisfied, not even for one particle in a one-dimensional grid. This can be verified forcing the above approximated expressions for \mathbf{J} and ρ into the discrete continuity equation:

$$\begin{aligned} \frac{\rho_{i,j,k}^{n+1} - \rho_{i,j,k}^n}{\Delta t} &+ \frac{J_{x_{i+1/2,j,k}}^{n+1/2} - J_{x_{i-1/2,j,k}}^{n+1/2}}{\Delta x} \\ &+ \frac{J_{y_{i,j+1/2,k}}^{n+1/2} - J_{y_{i,j-1/2,k}}^{n+1/2}}{\Delta y} \\ &+ \frac{J_{z_{i,j,k+1/2}}^{n+1/2} - J_{z_{i,j,k-1/2}}^{n+1/2}}{\Delta z} = 0, \end{aligned} \quad (4.2)$$

where a second order discretization is used for both space and time derivatives consistently with the chosen Maxwell solver. In practice, using the shape functions leads to discrepancies between the grid-defined charge and current fields, by construction. The resulting numerical error is of order $O(1)$, thus it cannot be controlled and a net charge may be deposited onto the grid during the simulations. For these reasons it is important to have a charge-conserving method. Nevertheless, the standard method has the important advantage of being energy-conserving, namely consistent with the single particle energy balance given by (1.15), hence it conserves the total energy of the system. Moreover, it is easy to implement and very efficient. Another type of deposition method exists that, on the contrary, is charge-conserving, but not energy-conserving. Different schemes have been proposed for example by Villasenor and Buneman [71], Esirkepov [70] and Umeda *et al.* [95]. Here the Esirkepov scheme is chosen because more efficient than the first, that requires many `if` statements, and easier to implement than the third when using second-order shape functions. To guarantee charge conservation it is also possible to use the non-conservative current and adjust the longitudinal part of the electric field \mathbf{E} by solving the Poisson equation (see [57] for details). A correction for \mathbf{E} is then introduced as \mathbf{E}' so that:

$$\operatorname{div} \mathbf{E}' = \rho, \quad \mathbf{E}' = \mathbf{E} - \nabla \delta\phi,$$

where

$$\Delta \delta\phi = \operatorname{div} \mathbf{E} - \rho. \quad (4.3)$$

This strategy is computationally inefficient if compared to those methods that directly build a conserving \mathbf{J} , because it requires to solve the Poisson problem (4.3) for each time step. Indeed, most of the available PIC codes implement one of the charge-conserving current deposition methods.

4.2.1 Esirkepov algorithm

Esirkepov's approach consists in computing the current density field from the continuity equation, thus straightly enforcing exact charge conservation [70]. It is also called *density decomposition scheme*. It works for Cartesian geometry and arbitrary shape functions, assuming that the macro-particles move along straight lines over one time step, without requiring to solve Poisson equation. The basic idea is to compute the local current density from the discrete continuity equation written for the single particle. The charge density field is again computed using the particles shape function as in formula (4.1), here reported in a more compact form:

$$\rho_{i,j,k}^n = \omega \sum_{p=1}^{N_p} \hat{B}_{i,j,k}(\mathbf{r}_p^n),$$

where \mathbf{r}_p^n is the position of the p -th particle at time t and $\omega = q/\Delta x \Delta y \Delta z$. Since the continuity equation (4.2) is linear, it is possible to first compute the current related to a

single particle, denoted by $(\mathcal{J}_x, \mathcal{J}_y, \mathcal{J}_z)$, and then sum up the contributions over all the particles. Let's define an auxiliary vector \mathbf{W} , called *density decomposition vector*, whose components are:

$$(W_x)_{i,j,k} = \frac{\Delta t}{\omega} \frac{\mathcal{J}_{x_{i+1/2,j,k}}^{n+1/2} - \mathcal{J}_{x_{i-1/2,j,k}}^{n+1/2}}{\Delta x} \quad (4.4)$$

$$(W_y)_{i,j,k} = \frac{\Delta t}{\omega} \frac{\mathcal{J}_{y_{i,j+1/2,k}}^{n+1/2} - \mathcal{J}_{y_{i,j-1/2,k}}^{n+1/2}}{\Delta y} \quad (4.5)$$

$$(W_z)_{i,j,k} = \frac{\Delta t}{\omega} \frac{\mathcal{J}_{z_{i,j,k+1/2}}^{n+1/2} - \mathcal{J}_{z_{i,j,k-1/2}}^{n+1/2}}{\Delta z}. \quad (4.6)$$

Summing together the components of \mathbf{W} , using the continuity equation and the definition of the discrete charge density field leads to (dropping grid indexes for simplicity):

$$W_x + W_y + W_z = \hat{B}(\mathbf{r}_p^{n+1}) - \hat{B}(\mathbf{r}_p^n).$$

Assume that \mathbf{W} is a linear combination of the following quantities that are related to a shift of the particle in the three-dimensional space:

$$\begin{aligned} & \hat{B}(x^n, y^n, z^n) \\ & \hat{B}(x^{n+1}, y^n, z^n) \\ & \hat{B}(x^n, y^{n+1}, z^n) \\ & \hat{B}(x^n, y^n, z^{n+1}) \\ & \hat{B}(x^{n+1}, y^{n+1}, z^n) \\ & \hat{B}(x^{n+1}, y^n, z^{n+1}) \\ & \hat{B}(x^n, y^{n+1}, z^{n+1}) \\ & \hat{B}(x^{n+1}, y^{n+1}, z^{n+1}) \end{aligned} \quad (4.7)$$

where $\mathbf{r}_p^n = (x^n, y^n, z^n)$ and $\mathbf{r}_p^{n+1} = (x^{n+1}, y^{n+1}, z^{n+1})$. Note that the density decomposition vector depends on the position of the particle at two subsequent times (n and $n+1$), as opposed to the standard method where \mathbf{J} is computed only from the position at time $n+1/2$ (see figure 4.6). From the former assumption and taking into account other properties that \mathbf{W} must satisfy (e.g. if the particle does not shift along one direction, say x , then the corresponding \mathbf{W} component is zero, $W_x = 0$), it can be demonstrated that only

one linear combination of the (4.7) is admissible:

$$\begin{aligned}
W_x &= \frac{1}{3}[\hat{B}(x^{n+1}, y^{n+1}, z^{n+1}) - \hat{B}(x^n, y^{n+1}, z^{n+1})] \\
&\quad + \frac{1}{6}[\hat{B}(x^{n+1}, y^n, z^{n+1}) - \hat{B}(x^n, y^n, z^{n+1})] \\
&\quad + \frac{1}{6}[\hat{B}(x^{n+1}, y^{n+1}, z^n) - \hat{B}(x^n, y^{n+1}, z^n)] \\
&\quad + \frac{1}{3}[\hat{B}(x^{n+1}, y^n, z^n) - \hat{B}(x^n, y^n, z^n)] \\
W_y &= \frac{1}{3}[\hat{B}(x^{n+1}, y^{n+1}, z^{n+1}) - \hat{B}(x^{n+1}, y^n, z^{n+1})] \\
&\quad + \frac{1}{6}[\hat{B}(x^n, y^{n+1}, z^{n+1}) - \hat{B}(x^n, y^n, z^{n+1})] \\
&\quad + \frac{1}{6}[\hat{B}(x^{n+1}, y^{n+1}, z^n) - \hat{B}(x^{n+1}, y^n, z^n)] \\
&\quad + \frac{1}{3}[\hat{B}(x^n, y^{n+1}, z^n) - \hat{B}(x^n, y^n, z^n)] \\
W_z &= \frac{1}{3}[\hat{B}(x^{n+1}, y^{n+1}, z^{n+1}) - \hat{B}(x^{n+1}, y^{n+1}, z^n)] \\
&\quad + \frac{1}{6}[\hat{B}(x^n, y^{n+1}, z^{n+1}) - \hat{B}(x^n, y^{n+1}, z^n)] \\
&\quad + \frac{1}{6}[\hat{B}(x^{n+1}, y^n, z^{n+1}) - \hat{B}(x^{n+1}, y^n, z^n)] \\
&\quad + \frac{1}{3}[\hat{B}(x^n, y^n, z^{n+1}) - \hat{B}(x^n, y^n, z^n)].
\end{aligned} \tag{4.8}$$

This formula is valid for an arbitrary shape function. From now on, second-order polynomial shape functions are used. Consider a particle located inside the cell centred in (i_1, j_1, k_1) at time step n that moves to the cell centred in (i_2, j_2, k_2) at time step $n + 1$. Since $\Delta t < \min\{\Delta x, \Delta y, \Delta z\}$, the particle can cross at most one border of the initial cell. This means that the final cell can only have $i_2 \in \{i_1 - 1, i_1, i_1 + 1\}$, $j_2 \in \{j_1 - 1, j_1, j_1 + 1\}$ and $k_2 \in \{k_1 - 1, k_1, k_1 + 1\}$. Hence the vector $\mathbf{W} = (W_x, W_y, W_z)$ has non-zero components only for $i \in \{i_1 - 2, i_1 - 1, i_1, i_1 + 1, i_1 + 2\}$, $j \in \{j_1 - 2, j_1 - 1, j_1, j_1 + 1, j_1 + 2\}$ and $k \in \{k_1 - 2, k_1 - 1, k_1, k_1 + 1, k_1 + 2\}$. The current related to a single particle is computed inverting (4.6) and taking into account that it is equal to zero on the grid nodes where \mathbf{W} is zero ("far" from the particle location). For example, the x component of the local current is:

$$\begin{aligned}
(\mathcal{J}_x)_{i_1+1/2+2,j,k} &= 0 \\
(\mathcal{J}_x)_{i_1+1/2+1,j,k} &= (\mathcal{J}_x)_{i_1+1/2+2,j,k} - \omega \frac{\Delta x}{\Delta t} (W_x)_{i_1+2,j,k} \\
(\mathcal{J}_x)_{i_1+1/2+0,j,k} &= (\mathcal{J}_x)_{i_1+1/2+1,j,k} - \omega \frac{\Delta x}{\Delta t} (W_x)_{i_1+1,j,k} \\
&\dots
\end{aligned}$$

The total current \mathbf{J} is computed summing all particles contributions $(\mathcal{J}_x, \mathcal{J}_y, \mathcal{J}_z)$ on every grid point. This procedure ensures the continuity equation and Gauss law are satisfied

within machine accuracy at every time step. Nevertheless, energy-conservation is not guaranteed anymore.

The Esirkepov method has the important drawback of being more expensive than the standard method. Using second-order shape functions for every particle, for every spatial dimension, 5 weights instead of 3 have to be computed. In addition, in a 3D simulation, once the 125×3 weights are known, the auxiliary vector containing the density decomposition must be calculated, while in the standard case one can straightly compute the current density from the 9×3 weights.

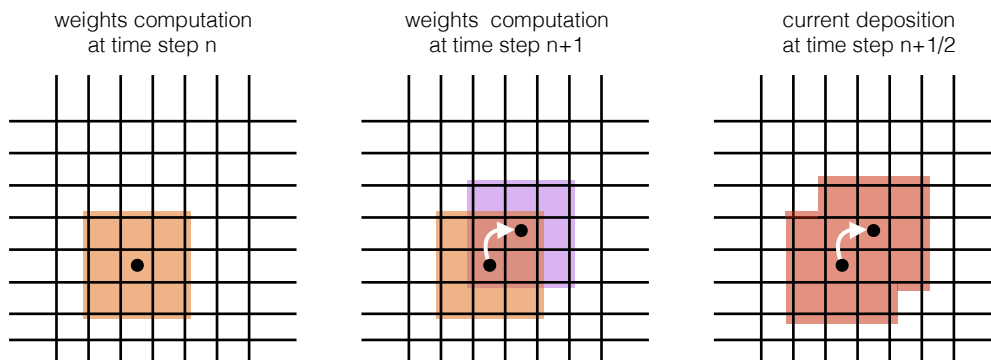


Figure 4.6: The Esirkepov method uses the particle weights at two subsequent time steps. These weights are used to compute the density decomposition vector, whose components are then used to compute the current density field.

The 1D and 2D versions of the algorithm require a simplified version of the formula 4.8. It is possible to obtain the 2D reduction imposing the current density homogeneous along the third dimension. Similarly, the 1D formula can be obtained from the 2D formula assuming the current to be homogeneous along one more dimension.

The pseudocode for the charge-conserving current deposition routine is the following.

4.2.2 Numerical tests

The density decomposition method has been implemented into the code *piccante*, which only came with an energy-conserving current deposition routine (called *standard* current deposition). A first set of simulations has been performed to test the 1D algorithm. Two different scenarios are considered. First, a box filled with a low-temperature plasma without any external sources. Second a partially filled box with a low-temperature plasma irradiated by a linearly polarized laser pulse, so to produce strongly nonlinear plasma waves. Using the Esirkepov method, the continuity equation and Gauss's law are satisfied almost within machine accuracy, as shown in figure 4.7. On the other hand, energy conservation is not satisfied as well as with the standard method (see figure 4.8). A secular increase in

```

for each particle do
  for each spatial dimension do
    find nearest integer and half-integer nodes and corresponding positions on
    the grid;
    compute particle weights;
    advance particle position;
    find new nearest integer and half-integer nodes and corresponding positions
    on the grid;
    compute new particle weights;
    compute difference between new and old weights;
  end
  for every node involved in the deposition do
    compute density decomposition vector;
    compute particle contribution to current density;
    add particle contribution to total current density;
  end
end

```

total energy due to the finite time step size, called *numerical heating*, is always observed (see Hockney, Eastwood [58]). Its effect is more important with the charge-conserving scheme, even though remains limited and seems to be acceptable.

Different tests were also performed using the 2D algorithm. Again, as a simple situation a box filled with a low-temperature plasma has been considered. As in the 1D case, Poisson equation and Gauss law are satisfied almost within machine accuracy and energy conservation is worse with the charge-conserving deposition, but still acceptable. As a more complex and interesting scenario the Esirkepov method has been used in the two-dimensional simulations presented in chapter 3, where a multi-layer foam-attached target is irradiated by a high intensity laser pulse. Two different configurations have been analysed varying the foam type: the uniform foam and one of the nanostructured foams made of regularly arranged balls (in particular foam 2, see figures 3.17 and 3.18). The differences on the results obtained with the two strategies for current deposition are qualitatively and quantitatively limited. Indeed, the main quantities such as electron density, electromagnetic field, proton energy spectra, proton distribution in angle and energy exhibit a very similar behaviour (as an example see figure 4.10). The main difference is found in the evolution of the total energy (see figure 4.9). With Esirkepov method energy conservation is "worse", but the increment is about 0.5 ‰, which is acceptable. For longer simulation times this increment is higher but remains limited (often errors $\lesssim 1\%$ are considered acceptable in PIC simulations). The main drawback is the doubling of the computational




Figure 4.7: Residual error: Gauss law (green) and Poisson equation (purple) for both the Esirkepv current deposition method (charge-conserving) and the standard scheme (energy-conserving). Note that the ordinate scales differ by about 14 orders of magnitude.




Figure 4.8: Total energy vs time for both the Esirkepv current deposition method (charge-conserving) and the standard scheme (energy-conserving).

time needed to perform a fixed amount of steps for every 2D simulation. Nevertheless, *piccante* performances with the Esirkepov method resulted comparable with the performances of the open-source particle-in-cell code Smilei [96], which uses Esirkepov method by default. Overall, the differences found with the two algorithms are limited, so that future simulations will use the standard algorithm because of its efficiency. In general, at the beginning of a numerical campaign, a good practice would be to test a given scenario with both algorithms and then, if very similar results are obtained, use the standard algorithm.




Figure 4.9: Evolution of total energy in time with the uniform foam (left) and with the nanostructured foam (right). In both cases energy conservation is "worse" using the Esirkepov algorithm, anyway the relative increase remains small, about 0.5 %.




Figure 4.10: Proton energy spectra with the uniform foam (left) and with the nanostructured foam (right). Similar results are obtained using the standard algorithm and the Esirkepov algorithm.

4.3 Conclusions

The developing work on particle-in-cell code *piccante* has been discussed. It consisted in the implementation of different strategies for adding OpenMP threading to the existent MPI code and a new charge-conserving current deposition algorithm.

The new OpenMP parallelization can be successfully exploited when the plasma is localized in one or more regions of the box, i.e. it is far from being uniformly distributed, which is not unusual for nanostructured plasmas. Anyway, as already mentioned, improvements need to be made to the present version of the "auxiliary currents" strategy for the OpenMP parallelization of the current deposition routine. The sizes of the copies of the current density vector shall be tuned depending on the number of threads. The hybridization of the code seems to lead to improvements in the performances in novel multi-core architectures, as suggested by few preliminary tests. Not only further investigations are required, but also additional optimizations are recommended to run the code on the new generation of Xeon Phi processors, such as parallelism on 512-bit SIMD level, i.e. vectorization.

The Esirkepov method allows to guarantee the continuity equation and Gauss law within machine accuracy, without solving Poisson equation. It demonstrated to run slower than the former current deposition algorithm and to ensure energy conservation with less precision, but still widely acceptable. In any case, the choice of the current deposition algorithm is problem-dependent.

Chapter 5

Conclusions and perspectives

Three-dimensional and two-dimensional kinetic particle-in-cell simulations have been performed as numerical investigations on laser-driven ion acceleration in an enhanced regime (see chapter 3). The system under examination consists in a high intensity laser pulse irradiated onto a double-layer foam-attached target (see section 1.5). The illuminated layer is a near-critical nanostructured foam, used to improve the acceleration performances through an enhanced target normal sheath acceleration-like process. This first layer is attached to a second layer, consisting in a conventional thin solid foil. Three-dimensional nanostructured foams are modelled through the diffusion-limited aggregation method, which produces numerical foams with quite realistic morphologies and structures (see section 3.1). Besides density non-uniformities, also (average) thickness non-uniformities have been considered to simulate what experimentally is an inherent consequence of the manufacturing process that may lead to different results for what concerns ion acceleration. On the other hand, two-dimensional nanostructured foams are simply modelled as collections of regularly arranged nanospheres with varying radii and centre-to-centre distance to simulate different nanostructures.

The particle-in-cell method proves to be an adequate tool for theoretical numerical investigations of the aforementioned systems, being efficient and flexible. The key features of the physics of the interaction, such as sheath formation and first stages of impurity proton acceleration, are well captured by simulations (see 3.4). Moreover, numerical results exhibit many features of the phenomena that are also observed in experiments, as the peculiar shapes of angular distribution of the emitted protons. The presence of the nanostructure significantly affects the interaction mechanisms, leading to an increase in laser energy absorption and, ultimately, in proton cutoff energy. Furthermore, the specific properties of the nanostructure play an important role in the enhancement degree of the acceleration process: different foam parameters lead to different acceleration performances. In particular, a dependence on the (mean) thickness of the region hit by the laser pulse was observed

with three-dimensional simulations (see paragraph 3.4.1); a dependence on the size and distance of the nanospheres was revealed by two-dimensional numerical studies (see paragraph 3.4.2). In any case it seems that, ideally, a uniform foam would lead to better results, though laboratory production of low-density uniform foams is non-trivial indeed. Simulation parameters have been chosen partially similar to those used in few past experimental works to allow reasonable comparisons; numerical output results showed good agreement with available experimental data. However, it is not possible to adopt a completely realistic simulation setup, not even if a fully-optimized and parallel code runs on state-of-the-art supercomputers (see section 3.5). The allowed setups are limited due to the high computational resources required by this kind of simulations, especially in three dimensions. Hence, a further optimization of the PIC code *piccante*, used to perform all simulations, resulted convenient. The OpenMP parallelization model has been developed into the code, showing good performance results for very special plasma configurations (see section 4.1). Satisfactory results have been obtained when running the code on Intel MIC coprocessors, promising to achieve even higher performances on next-generation MIC architectures, soon to be delivered at CINECA. Furthermore, the Esirkepov current deposition method has been added into *piccante* to understand whether exact charge-conservation is mandatory or not for what concerns the systems under investigation (see section 4.2). It was revealed that, at least in those situations, the more efficient but not charge-conserving standard algorithm gives satisfactory outcomes, being the results obtained with the two methods quite similar.

Several issues naturally arise as a continuation of the presented work. They are summarized in the following points.

- *Mathematical foam models.* Reliable modelling of the nanostructured foam is advisable, therefore the specific growth process should be carefully taken into account through suitable mathematical models. Various extensions of the DLA model could be considered so to compare the different results with experimental observations on the growth regime. Also simpler morphologies could result interesting as for example regularly or randomly arranged nanowires (instead of nanospheres). Besides growth mechanism, also foam properties should be taken into account. Material composition, size of the building blocks (e.g. nanoparticles), average density, average thickness always need to be considered with parametric scans when searching the optimal configuration.
- *Further code developments in *piccante*.* Additional code refinements may result interesting both to improve code performances and include new physical features. Further optimizations may concern the implementation of parallelism on 512-bit SIMD level, i.e. vectorization, which could decrease the current CPU time needed to perform a given simulation by a factor of four in new supercomputer architectures. Moreover,

additional physics could be introduced such as ionization, collisions (even if negligible play an important role in transport phenomena), quantum effects and a Poisson solver (to allow a non globally neutral initial condition).

- *Simulations of laser-driven ion acceleration experiments.* Multi-dimensional numerical campaigns could be held in order to analyse the role and the behaviour of different experimental configurations, varying not only the target properties (as already mentioned), but also the laser pulse properties such as polarization, duration, intensity, contrast, angle of incidence. These studies are useful to be able to propose new solutions to improve the performance of the acceleration scheme, giving ideas for novel target concepts and experimental campaigns.

Based on these ideas, the present work may represent a starting point for future, deeper numerical studies on laser-driven ion acceleration with nanostructured materials. PIC simulations may be employed in the future to understand new observations, design novel experimental concepts and support laboratory design of the targets. Moreover, they represent the only available tool to envisage what occurs with parameters currently beyond reach, e.g. at intensities $> 10^{22}$ W/cm². New experimental configurations may be considered, such as different foam models, foam characteristics and laser properties. Extensive scans may be performed with two-dimensional simulations varying one of these parameters at a time; then, the resulting optimum configuration may be simulated in three dimensions to deeply investigate the whole system. Such investigations would likely allow to get a more satisfactory understanding of the physics of interaction between high intensity laser pulses and nanostructured materials.

Appendix A

Klimontovich approach

Here a simple procedure to derive Klimontovich equation from the equations of motion of the plasma particles is reported. Let the indexes i, a denote the quantities related to the i -th particle of kind a , with mass m_a and charge q_a . The relativistic dynamic of each particle constituting the plasma is the solution of the following system ($i = 1, \dots, N_a; a = 1, \dots, N$):

$$\begin{cases} \frac{d\mathbf{r}_{i,a}(t)}{dt} = \frac{\mathbf{p}_{i,a}(t)}{m_a \gamma_{i,a}(t)} \\ \frac{d\mathbf{p}_{i,a}(t)}{dt} = q_a \left[\mathbf{E}_{micr}(\mathbf{r}_{i,a}(t), t) + \frac{\mathbf{p}_{i,a}(t)}{m_a c \gamma_{i,a}(t)} \times \mathbf{B}_{micr}(\mathbf{r}_{i,a}(t), t) \right], \end{cases} \quad (\text{A.1})$$

where $\mathbf{r}_{i,a}$ and $\mathbf{p}_{i,a}$ are the particle position and momenta respectively, while $\gamma_{i,a}$ is its relativistic factor defined by

$$\gamma_{i,a}(t) = \sqrt{1 + \frac{\mathbf{p}_{i,a}(t) \cdot \mathbf{p}_{i,a}(t)}{m_a^2 c^2}}.$$

The particle velocity $\mathbf{v}_{i,a}$ is related to the momenta by the following relation

$$\mathbf{p}_{i,a}(t) = m_a \gamma_{i,a}(t) \mathbf{v}_{i,a}(t).$$

The microscopic electromagnetic field $\mathbf{E}_{micr} = \mathbf{E}_{micr}(\mathbf{r}, t)$, $\mathbf{B}_{micr} = \mathbf{B}_{micr}(\mathbf{r}, t)$ is the solution of the Maxwell system:

$$\left\{ \begin{array}{l} \text{div } \mathbf{E}_{micr}(\mathbf{r}, t) = 4\pi \rho_{micr}^{tot}(\mathbf{r}, t) \\ \text{curl } \mathbf{B}_{micr}(\mathbf{r}, t) = \frac{4\pi}{c} \mathbf{J}_{micr}^{tot}(\mathbf{r}, t) + \frac{1}{c} \frac{\partial \mathbf{E}_{micr}(\mathbf{r}, t)}{\partial t} \\ \text{div } \mathbf{B}_{micr}(\mathbf{r}, t) = 0 \\ \text{curl } \mathbf{E}_{micr}(\mathbf{r}, t) = -\frac{1}{c} \frac{\partial \mathbf{B}_{micr}(\mathbf{r}, t)}{\partial t}, \end{array} \right. \quad (\text{A.2})$$

where the sources for the electromagnetic field ρ_{micr}^{tot} e \mathbf{J}_{micr}^{tot} comprise both the contribution of the external charges and plasma charged populations:

$$\begin{aligned} \rho_{micr}^{tot} &= \rho_{ext} + \sum_{a=1}^N \rho_{micr,a} \\ \mathbf{J}_{micr}^{tot} &= \mathbf{J}_{ext} + \sum_{a=1}^N \mathbf{J}_{micr,a}, \end{aligned}$$

where

$$\begin{aligned} \rho_{micr,a}(\mathbf{r}, t) &= q_a \sum_{i=1}^{N_a} \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) \\ \mathbf{J}_{micr,a}(\mathbf{r}, t) &= q_a \sum_{i=1}^{N_a} \mathbf{v}_{i,a}(t) \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)). \end{aligned}$$

Consider the six-dimensional phase space of Eulerian coordinates $(\mathbf{r}, \mathbf{p}_a)$, position and a -th species momenta respectively, where $\mathbf{p}_a = m_a \gamma_a \mathbf{v}$ and $\gamma_a = \sqrt{1 + \frac{\mathbf{p}_a \cdot \mathbf{p}_a}{m_a^2 c^2}}$, \mathbf{v} being the Eulerian velocity. So far, the description is partially Lagrangian (particles motion) and partially Eulerian (electromagnetic field). To write the whole problem using an Eulerian approach, the microscopic distribution function is introduced as:

$$f_{micr,a}(\mathbf{r}, \mathbf{p}_a, t) = \sum_{i=1}^{N_a} \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) \delta(\mathbf{p}_a - \mathbf{p}_{i,a}(t)),$$

where $\mathbf{r}_{i,a}$ e $\mathbf{p}_{i,a}$ are the solutions to the differential system (A.1). By definition $f_{micr,a}(\mathbf{r}, \mathbf{p}_a, t) d\mathbf{r} d\mathbf{p}_a$ is equal to the number of particles of kind a whose positions and momenta at time t lie around the point $(\mathbf{r}, \mathbf{p}_a)$ of phase space in the range $d\mathbf{r} d\mathbf{p}_a$. The microscopic sources can

be expressed in terms of f_a :

$$\begin{aligned}\rho_{micr,a}^{tot} &= q_a \sum_{i=1}^{N_a} \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) = q_a \int f_{micr,a}(\mathbf{r}, \mathbf{p}_a, t) d\mathbf{p}_a \\ \mathbf{J}_{micr,a}^{tot} &= q_a \sum_{i=1}^{N_a} \mathbf{v}_{i,a}(t) \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) = q_a \int \mathbf{v} f_{micr,a}(\mathbf{r}, \mathbf{p}_a, t) d\mathbf{p}_a.\end{aligned}$$

Klimontovich equation, which is a sort of continuity equation in phase space, can be deduced by simply computing the partial time derivative of $f_{micr,a}$ and using some properties of the Dirac delta functions. It holds:

$$\begin{aligned}\frac{\partial f_{micr,a}}{\partial t} &= \sum_{i=1}^{N_a} \frac{\partial}{\partial t} [\delta(\mathbf{r} - \mathbf{r}_{i,a}) \delta(\mathbf{p}_a - \mathbf{p}_{i,a})] \\ &= \sum_{i=1}^{N_a} \left[\delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \frac{\partial}{\partial t} \delta(\mathbf{r} - \mathbf{r}_{i,a}) + \delta(\mathbf{r} - \mathbf{r}_{i,a}) \frac{\partial}{\partial t} \delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \right] \\ &= \sum_{i=1}^{N_a} \left[\delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \dot{\mathbf{r}}_{i,a} \cdot \frac{\partial}{\partial \mathbf{r}_{i,a}} \delta(\mathbf{r} - \mathbf{r}_{i,a}) + \delta(\mathbf{r} - \mathbf{r}_{i,a}) \dot{\mathbf{p}}_{i,a} \cdot \frac{\partial}{\partial \mathbf{p}_{i,a}} \delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \right].\end{aligned}$$

Dirac delta functions satisfy the following property:

$$\frac{\partial}{\partial \mathbf{a}} \delta(\mathbf{a} - \mathbf{b}) = -\frac{\partial}{\partial \mathbf{b}} \delta(\mathbf{a} - \mathbf{b}),$$

hence

$$\begin{aligned}\frac{\partial f_{micr,a}}{\partial t} &= -\sum_{i=1}^{N_a} \left[\delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \dot{\mathbf{r}}_{i,a} \cdot \frac{\partial}{\partial \mathbf{r}} \delta(\mathbf{r} - \mathbf{r}_{i,a}) + \delta(\mathbf{r} - \mathbf{r}_{i,a}) \dot{\mathbf{p}}_{i,a} \cdot \frac{\partial}{\partial \mathbf{p}_a} \delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \right] \\ &= -\operatorname{div}_{\mathbf{r}} \left(\sum_{i=1}^{N_a} \dot{\mathbf{r}}_{i,a} \delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \delta(\mathbf{r} - \mathbf{r}_{i,a}) \right) - \operatorname{div}_{\mathbf{p}_a} \left(\sum_{i=1}^{N_a} \dot{\mathbf{p}}_{i,a} \delta(\mathbf{r} - \mathbf{r}_{i,a}) \delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \right).\end{aligned}$$

Now, the following can be applied

$$\delta(\mathbf{a} - \mathbf{b}) f(\mathbf{a}) = \delta(\mathbf{a} - \mathbf{b}) f(\mathbf{b}),$$

so that

$$\begin{aligned}\frac{\partial f_{micr,a}}{\partial t} &= -\operatorname{div}_{\mathbf{r}} \left(\dot{\mathbf{r}} \sum_{i=1}^{N_a} \delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \delta(\mathbf{r} - \mathbf{r}_{i,a}) \right) - \operatorname{div}_{\mathbf{p}_a} \left(\dot{\mathbf{p}}_a \sum_{i=1}^{N_a} \delta(\mathbf{r} - \mathbf{r}_{i,a}) \delta(\mathbf{p}_a - \mathbf{p}_{i,a}) \right) \\ &= -\operatorname{div}_{\mathbf{r}} (\dot{\mathbf{r}} f_{micr,a}) - \operatorname{div}_{\mathbf{p}_a} (\dot{\mathbf{p}}_a f_{micr,a}).\end{aligned}$$

Klimontovich equation for the function $f_a = f_a(\mathbf{r}, \mathbf{p}_a, t)$ is then recovered:

$$\frac{\partial f_{micr,a}}{\partial t} + \operatorname{div}_{\mathbf{r}} \left(\frac{\mathbf{p}_a}{m_a c \gamma_a} f_{micr,a} \right) + \operatorname{div}_{\mathbf{p}_a} \left[q_a \left(\mathbf{E}_{micr}(\mathbf{r}, t) + \frac{\mathbf{p}_a}{m_a c \gamma_a} \times \mathbf{B}_{micr}(\mathbf{r}, t) \right) f_{micr,a} \right] = 0 \quad (\text{A.3})$$

Because of the nature of the Lorentz force equation A.3 can be written as :

$$\left[\frac{\partial}{\partial t} + \frac{\mathbf{p}_a}{m_a \gamma_a} \cdot \nabla_{\mathbf{r}} + q_a \left(\mathbf{E}_{micr}(\mathbf{r}, t) + \frac{\mathbf{p}_a}{m_a c \gamma_a} \times \mathbf{B}_{micr}(\mathbf{r}, t) \right) \cdot \nabla_{\mathbf{p}_a} \right] f_{micr,a} = 0,$$

where was used the formula

$$\operatorname{div}(\mathbf{a} \times \mathbf{b}) = \operatorname{curl} \mathbf{a} \cdot \mathbf{b} - \operatorname{curl} \mathbf{b} \cdot \mathbf{a}$$

and the fact that

$$\operatorname{curl}_{\mathbf{p}_a} \left(\frac{\mathbf{p}_a}{m_a c \gamma_a} \right) = 0. \quad (\text{A.4})$$

This proves that the equations of motion imply Klimontovich equation. The opposite implication can be proven by substituting into Klimontovich equation the expression of the microscopic distribution function. Defining

$$\mathbf{F}_{micr}(\mathbf{r}, \mathbf{p}_a, t) = q_a \left(\mathbf{E}_{micr}(\mathbf{r}, t) + \frac{\mathbf{p}_a}{m_a c \gamma_a} \times \mathbf{B}_{micr}(\mathbf{r}, t) \right),$$

it yields:

$$\left[\frac{\partial}{\partial t} + \frac{\mathbf{p}_a}{m_a \gamma_a} \cdot \nabla_{\mathbf{r}} + \mathbf{F}_{micr} \cdot \nabla_{\mathbf{p}_a} \right] \sum_{i=1}^{N_a} \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) \delta(\mathbf{p}_a - \mathbf{p}_{i,a}(t)) = 0.$$

This can be written as:

$$\sum_{i=1}^{N_a} \left[-\delta(\mathbf{p} - \mathbf{p}_{i,a}(t)) \dot{\mathbf{r}}_{i,a} \cdot \nabla \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) - \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) \dot{\mathbf{p}}_{i,a} \cdot \nabla \delta(\mathbf{p}_a - \mathbf{p}_{i,a}(t)) \right. \\ \left. + \delta(\mathbf{p} - \mathbf{p}_{i,a}(t)) \frac{\mathbf{p}_a}{m_a \gamma_a} \cdot \nabla \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) + \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) \mathbf{F}_{micr}(\mathbf{r}, \mathbf{p}_a, t) \cdot \nabla \delta(\mathbf{p}_a - \mathbf{p}_{i,a}(t)) \right] = 0.$$

Integrating over momenta, the equation for the evolution of the particle position is recovered

$$\sum_{i=1}^{N_a} \left[-\frac{d\mathbf{r}_{i,a}(t)}{dt} + \frac{\mathbf{p}_{i,a}(t)}{m_a \gamma_{i,a}(t)} \right] \cdot \nabla \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) = 0 \quad \forall \mathbf{r}, \forall t \\ \implies \frac{d\mathbf{r}_{i,a}(t)}{dt} = \frac{\mathbf{p}_{i,a}(t)}{m_a \gamma_{i,a}(t)} \quad \forall i, \forall t, \quad (\text{A.5})$$

where it was used the following relation

$$\int \mathbf{F}_{micr}(\mathbf{r}, \mathbf{p}_a, t) \cdot \nabla \delta(\mathbf{p}_a - \mathbf{p}_{i,a}(t)) d\mathbf{p}_a = - \int \delta(\mathbf{p}_a - \mathbf{p}_{i,a}(t)) \operatorname{div}_{\mathbf{p}_a} \mathbf{F}_{micr}(\mathbf{r}, \mathbf{p}_a, t) d\mathbf{p}_a = 0.$$

Integrating over positions leads to

$$\delta(\mathbf{p}_a - \mathbf{p}_{i,a}(t)) \int \left(-\dot{\mathbf{r}}_{i,a} + \frac{\mathbf{p}_a}{m_a \gamma_a} \right) \cdot \nabla \delta(\mathbf{r} - \mathbf{r}_{i,a}(t)) d\mathbf{r} + \\ \nabla \delta(\mathbf{p}_a - \mathbf{p}_{i,a}(t)) \cdot (-\dot{\mathbf{p}}_{i,a} + \mathbf{F}_{micr}(\mathbf{r}_{i,a}(t), \mathbf{p}_a, t)) = 0.$$

Using the fact that $\operatorname{div}_{\mathbf{p}_a} \mathbf{F}_{micr} = 0$ (it follows from (A.4)), the equation for the evolution of the particles momenta is obtained:

$$\frac{d\mathbf{p}_{i,a}(t)}{dt} = \mathbf{F}_{micr}(\mathbf{r}_{i,a}(t), \mathbf{p}_{i,a}(t), t).$$

List of Figures

1.1	Principle of chirped pulse amplification.	4
1.2	Approximate magnitudes in some typical plasmas. Reprinted from the NRL Plasma Formulary. Here "laser plasmas" refers to the interaction with long pulses (ns durations) and relatively small intensities (10^{14} W/cm ²).	7
1.3	Example of phase space filamentation. Reprinted from [9].	12
1.4	Schematic action of the ponderomotive force.	17
1.5	Schematic wakefield acceleration. Taken from [19].	17
1.6	Artistic depiction of a typical laser-driven ion acceleration experiment. From [20].	18
1.7	Experimental scaling of the maximum proton energy as function of the pulse power. Data refer to different experimental campaigns performed in the facilities according to the legend. Left: ultra-short pulses with durations around 30 fs. Right: "long" pulses with durations of hundreds of fs; the big red dot represents the current record.	19
1.8	Example of the depth dependence of the relative dose, i.e. the energy absorbed per unit mass, for different radiations in water. Proton and C ion profiles are characterized by the Bragg peak at the end of the path.	20
1.9	Example of RCF stack obtained with the TRIDENT short-pulse laser at the Los Alamos National Laboratory. Indicated beneath each film is the proton Bragg peak energy. Reprinted from [28].	21
1.10	Schematic target normal sheath acceleration. The laser-target interaction produces hot electrons at the front side that move across the bulk and reach the rear side where they form a sheath in vacuum. The electric sheath field generated by the charge separation is almost perpendicular to the rear surface and is able to accelerate ions.	23
1.11	Schematic sketch of the targets. Conventional targets are overdense, μm -thick solid foils (left panel). Foam-attached targets are special double-layer targets where a near-critical foam layer is placed in front of a μm -thick solid foil to enhance the acceleration process (right panel).	24

1.12	Scanning Electron Microscopy images that show how the morphology of the foam targets depend on the ambient gas properties. Top views and cross sections alternate on the rows.	25
1.13	High Resolution Scanning Transmission Electron Microscopy image of a foam sample. The building blocks of the aggregate are nanoparticles of $\sim 10\text{-}20$ nm size. Reprinted from [37].	26
1.14	Scanning Electron Microscopy images showing the aggregates assembled on the micrometer scale. Reprinted from [37].	26
1.15	Energy spectra of the protons obtained with a laser intensity of 4.1, 3.5, 3.7×10^{20} W/cm ² for S, P and C polarization respectively. The targets are double-layer targets with 8 μm -thick foam (spectra a, b, c respectively) and single targets (spectra d, e and f). The spectra are collected along the target normal direction. The inset shows the electron energy spectra when using double-layer targets with 12 μm -thick foam with S, P and C polarization, (spectra a, b, c respectively) and using single targets for S and C polarization (spectra d and f). Reprinted from [33].	26
2.1	Sketch of the idea underlying semi-Lagrangian methods. Backwards (left), forward (middle) and finite volume (right). Taken from [43].	31
2.2	One time step in a particle-in-cell simulation program. Adapted from [57].	32
2.3	Sampling of the distribution function with macro-particles. Each macro-particle has a definite momentum, but is extended in space.	34
2.4	Different plasma models. Dashed arrows are transitions showing non-trivial effects. Reprinted from [63].	36
2.5	Splines of different order used as convolution kernels to couple particles with fields. Left: order 0, nearest grid point method. Middle: order 1, cloud-in-cell method. Right: order 2.	40
2.6	Yee-lattice. Reproduced from Wikipedia.	42
2.7	Main time cycle in a PIC code using standard algorithms.	45
3.1	Simplified scheme of a numerical multi-layer target: foam (dark blue), solid foil (red) and contaminant layer (light blue). Typical values of the electron density and thickness are indicated.	52
3.2	Graphical representation of a 3D DLA numerical foam.	53
3.3	Numerical foams obtained with the DLA method, with the hollow (left) and with the hill (right). Particles' colour follows the x coordinates.	59
3.4	Snapshot of the 3D simulation with the bare solid target ($t = 80$ fs).	61
3.5	Snapshot of the 3D simulation with the uniform foam-attached target ($t = 80$ fs).	61

3.6	Snapshot of the 3D simulation with the hollow foam-attached target ($t = 80$ fs).	62
3.7	Snapshot of the 3D simulation with the hill foam-attached target ($t = 80$ fs).	62
3.8	Electron density and electromagnetic energy density at plane $z = 0$.	63
3.9	Electron density and electromagnetic energy density at plane $z = 0$.	64
3.10	Electron density and electromagnetic energy density at plane $z = 0$.	65
3.11	Electron density and electromagnetic energy density at plane $z = 0$.	66
3.12	Time evolution of the energy normalized to the initial laser energy of foam and substrate electrons (left) and contaminant protons (right).	67
3.13	Contaminant protons energy spectra at time $t = 160$ fs, filtered in angles $\theta \in (-5, 5)$, $\phi \in (-5, 5)$.	68
3.14	Angular distribution of contaminant protons with energy $E \in [18, 22]$ MeV at the final time $t = 160$ fs. ϕ is the angle between the ion position and the target normal on the (x, y) plane, θ is the angle between the projection of the ion position on the (x, y) plane and z axis.	69
3.15	Angular distribution of contaminant protons with energy $E \in [0.9, 1.1]$ MeV. ϕ is the angle between the ion position and the target normal on the (x, y) plane, θ is the angle between the projection of the ion position on the (x, y) plane and z axis.	70
3.16	RCFs for materials with increasing conductivity from left to right. From [87].	73
3.17	Snapshots of the electron density of the foam-attached targets at time 27 fs (top) and 80 fs (bottom).	74
3.18	Snapshots of the electron density of the foam-attached targets at time 133 fs (top) and 187 fs (bottom).	75
3.19	Snapshots of the magnetic field B_z at time 27 fs (top) and 80 fs (bottom).	76
3.20	Snapshots of the magnetic field B_z at time 133 fs (top) and 187 fs (bottom).	77
3.21	Evolution of the energy repartition among electromagnetic field, electrons and ions.	78
3.22	Time evolution of the energy normalized to the initial laser energy of foam and substrate electrons (left) and contaminant protons (right).	78
3.23	Contaminant protons energy spectra at the end of the simulation filtered in angle $\phi \in (-2, 2)$.	79
3.24	Distribution in angle and energy of the contaminant protons at two time steps.	80
4.1	Scheme of the algorithm used to rearrange the array containing particles' data. All remaining particles need to be stored in its front. Starting from the back, as soon as a remaining particle is found, its data is copied in the position corresponding to the first leaving particle, and so on.	90
4.2	Partially filled box vs. completely filled box.	92

4.3	Partially filled box vs. completely filled box.	94
4.4	Snapshot of a simulation with a strongly non-uniform plasma, localized mainly along one direction (red axis). OpenMP parallelization along that direction together with MPI parallelization along the other two coordinates (yellow and green axes) may result convenient.	95
4.5	Initial condition with strongly non-uniform plasma for OpenMP performance testing. An hybrid execution showed a 1.5 gain in the elapsed time with respect to the purely MPI case.	95
4.6	The Esirkepov method uses the particle weights at two subsequent time steps. These weights are used to compute the density decomposition vector, whose components are then used to compute the current density field. . . .	100
4.7	Residual error: Gauss law (green) and Poisson equation (purple) for both the Esirkepov current deposition method (charge-conserving) and the standard scheme (energy-conserving). Note that the ordinate scales differ by about 14 orders of magnitude.	102
4.8	Total energy vs time for both the Esirkepov current deposition method (charge-conserving) and the standard scheme (energy-conserving).	102
4.9	Evolution of total energy in time with the uniform foam (left) and with the nanostructured foam (right). In both cases energy conservation is "worse" using the Esirkepov algorithm, anyway the relative increase remains small, about 0.5 ‰.	103
4.10	Proton energy spectra with the uniform foam (left) and with the nanostructured foam (right). Similar results are obtained using the standard algorithm and the Esirkepov algorithm.	103

List of Tables

3.1	Summary of three-dimensional simulations setup.	60
3.2	Centre-to-centre distance and radii of the spheres of the nanostructured foams.	72
3.3	Summary of two-dimensional simulations setup.	73
4.1	Summary of the simulations setup for the evaluation of code performances with OpenMP parallelization.	92
4.2	Performance results for partially filled box test.	93
4.3	Performance results for completely filled box test.	93

Bibliography

- [1] VI Veksler. The principle of coherent acceleration of charged particles. *The Soviet Journal of Atomic Energy*, 2(5):525–528, 1957.
- [2] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, C. B. Schroeder, Cs. Tóth, J. Daniels, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, and E. Esarey. Multi-GeV electron beams from capillary-discharge-guided subpetawatt laser pulses in the self-trapping regime. *Phys. Rev. Lett.*, 113:245002, Dec 2014.
- [3] F Wagner, O Deppert, C Brabetz, P Fiala, A Kleinschmidt, P Poth, VA Schanz, A Tebartz, B Zielbauer, M Roth, et al. Maximum proton energy above 85 MeV from the relativistic interaction of laser pulses with micrometer thick CH 2 targets. *Physical review letters*, 116(20):205002, 2016.
- [4] Milton Stanley Livingston. High energy accelerators. 1954.
- [5] Orazio Svelto and David C Hanna. *Principles of lasers*. Springer, 1976.
- [6] Donna Strickland and Gerard Mourou. Compression of amplified chirped optical pulses. *Optics communications*, 56(3):219–221, 1985.
- [7] Colin Danson, David Hillier, Nicholas Hopps, and David Neely. Petawatt class lasers worldwide. *High Power Laser Science and Engineering*, 3:e3, 2015.
- [8] Lewi Tonks and Irving Langmuir. Oscillations in ionized gases. *Phys. Rev.*, 33:195–210, Feb 1929.
- [9] Francis Filbet and Eric Sonnendrücker. Comparison of eulerian vlasov solvers. *Computer Physics Communications*, 150(3):247–266, 2003.
- [10] S.R. Groot, W.A. Leeuwen, and C.G. van Weert. *Relativistic kinetic theory: principles and applications*. North-Holland Pub. Co., 1980.
- [11] John David Jackson. *Classical electrodynamics*. Wiley, 1999.
- [12] A Pukhov and J Meyer-ter Vehn. Relativistic magnetic self-channeling of light in near-critical plasma: three-dimensional particle-in-cell simulation. *Physical review letters*, 76(21):3975, 1996.
- [13] M Borghesi, AJ MacKinnon, L Barringer, R Gaillard, LA Gizzi, C Meyer, O Willi, A Pukhov, and J Meyer-ter Vehn. Relativistic channeling of a picosecond laser pulse in a near-critical preformed plasma. *Physical review letters*, 78(5):879, 1997.

- [14] Paul Gibbon. *Short pulse laser interactions with matter*. World Scientific Publishing Company, 2004.
- [15] Andrea Macchi. *A Superintense Laser-Plasma Interaction Theory Primer*. Springer Science & Business Media, 2013.
- [16] Peter Mulser and Dieter Bauer. *High power laser-matter interaction*, volume 238. Springer Science & Business Media, 2010.
- [17] D Bauer, P Mulser, and W-H Steeb. Relativistic ponderomotive force, uphill acceleration, and transition to chaos. *Physical review letters*, 75(25):4622, 1995.
- [18] T Tajima and JM Dawson. Laser electron accelerator. *Physical Review Letters*, 43(4):267, 1979.
- [19] Thomas Katsouleas. Accelerator physics: Electrons hang ten on laser wake. *Nature*, 431(7008):515–516, 2004.
- [20] A. Macchi, M. Borghesi, and M. Passoni. Ion acceleration by superintense laser-plasma interaction. *Rev. Mod. Phys.*, 85:751–793, May 2013.
- [21] E. L. Clark, K. Krushelnick, J. R. Davies, M. Zepf, M. Tatarakis, F. N. Beg, A. Machacek, P. A. Norreys, M. I. K. Santala, I. Watts, and A. E. Dangor. Measurements of energetic proton transport through magnetized plasma from intense laser interactions with solids. *Phys. Rev. Lett.*, 84:670–673, Jan 2000.
- [22] A. Maksimchuk, S. Gu, K. Flippo, D. Umstadter, and V. Yu. Bychenkov. Forward ion acceleration in thin films driven by a high-intensity laser. *Phys. Rev. Lett.*, 84:4108–4111, May 2000.
- [23] RA Snavely, MH Key, SP Hatchett, TE Cowan, M Roth, TW Phillips, MA Stoyer, EA Henry, TC Sangster, MS Singh, et al. Intense high-energy proton beams from petawatt-laser irradiation of solids. *Physical Review Letters*, 85(14):2945, 2000.
- [24] Hiroyuki Daido, Mamiko Nishiuchi, and Alexander S Pirozhkov. Review of laser-driven ion sources and their applications. *Reports on Progress in Physics*, 75(5):056401, 2012.
- [25] S.V Bulanov, T.Zh Esirkepov, V.S Khoroshkov, A.V Kuznetsov, and F Pegoraro. Oncological hadrontherapy with laser ion accelerators. *Physics Letters A*, 299(2–3):240 – 247, 2002.
- [26] Ernesto Amato, Antonio Italiano, Daniele Margarone, Benedetta Pagano, Sergio Baldari, and Georg Korn. Study of the production yields of ^{18}F , ^{11}C , ^{13}N and ^{15}O positron emitters from plasma-laser proton sources at eli-beamlines for labeling of {PET} radiopharmaceuticals. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 811:1 – 5, 2016.
- [27] DP Higginson, L Vassura, MM Gugiu, Patrizio Antici, M Borghesi, S Brauckmann, C Diouf, A Green, L Palumbo, H Petrascu, et al. Temporal narrowing of neutrons produced by high-intensity short-pulse lasers. *Physical review letters*, 115(5):054802, 2015.
- [28] F Nürnberg, M Schollmeier, E Brambrink, A Blažević, DC Carroll, K Flippo, DC Gautier, M Geissel, K Harres, BM Hegelich, et al. Radiochromic film imaging spectroscopy of laser-accelerated proton beams. *Review of scientific instruments*, 80(3):033301, 2009.

- [29] SC Wilks, AB Langdon, TE Cowan, M Roth, M Singh, S Hatchett, MH Key, D Pennington, A MacKinnon, and RA Snavely. Energetic proton generation in ultra-intense laser–solid interactions. *Physics of Plasmas (1994-present)*, 8(2):542–549, 2001.
- [30] Claudio Perego, Alessandro Zani, Dimitri Batani, and Matteo Passoni. Extensive comparison among target normal sheath acceleration theoretical models. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 653(1):89–93, 2011.
- [31] Andrea Sgattoni, Pasquale Londrillo, Andrea Macchi, and Matteo Passoni. Laser ion acceleration using a solid target coupled with a low-density layer. *Physical Review E*, 85(3):036405, 2012.
- [32] Matteo Passoni, Alessandro Zani, Andrea Sgattoni, David Dellasega, Andrea Macchi, Irene Prencipe, Vincent Floquet, Philippe Martin, TV Liseykina, and Tiberio Ceccotti. Energetic ions at moderate laser intensities using foam-based multi-layered targets. *Plasma Physics and Controlled Fusion*, 56(4):045001, 2014.
- [33] M Passoni, A Sgattoni, I Prencipe, L Fedeli, D Dellasega, L Cialfi, Il Woo Choi, I Jong Kim, KA Janulewicz, Hwang Woon Lee, et al. Toward high-energy laser-driven ion beams: Nanostructured double-layer targets. *Physical Review Accelerators and Beams*, 19(6):061301, 2016.
- [34] JH Bin, WJ Ma, HY Wang, MJV Streeter, C Kreuzer, D Kiefer, M Yeung, S Cousens, PS Foster, B Dromey, et al. Ion acceleration using relativistic pulse shaping in near-critical-density plasmas. *Physical review letters*, 115(6):064801, 2015.
- [35] Andreas Henig, S Steinke, M Schnürer, T Sokollik, Rainer Hörlein, D Kiefer, D Jung, Jörg Schreiber, BM Hegelich, XQ Yan, et al. Radiation-pressure acceleration of ion beams driven by circularly polarized laser pulses. *Physical Review Letters*, 103(24):245003, 2009.
- [36] I Jong Kim, Ki Hong Pae, Chul Min Kim, Hyung Taek Kim, Jae Hee Sung, Seong Ku Lee, Tae Jun Yu, Il Woo Choi, Chang-Lyoul Lee, Kee Hwan Nam, et al. Transition of proton energy scaling using an ultrathin target irradiated by linearly polarized femtosecond laser pulses. *Physical review letters*, 111(16):165003, 2013.
- [37] A. Zani, D. Dellasega, V. Russo, and M. Passoni. Ultra-low density carbon foams produced by pulsed laser deposition. *Carbon*, 56:358–365, 2013.
- [38] Eric Sonnendrücker, Jean Roche, Pierre Bertrand, and Alain Ghizzo. The semi-lagrangian method for the numerical resolution of the Vlasov equation. *Journal of computational physics*, 149(2):201–220, 1999.
- [39] C.Z. Cheng and G. Knorr. The integration of the Vlasov equation in configuration space. *Journal of Computational Physics*, 22(3):330–351, 1976.
- [40] C.Z. Cheng. The integration of the Vlasov equation for a magnetized plasma. *Journal of Computational Physics*, 24(4):348–360, 1977.
- [41] T. Nakamura and T. Yabe. Cubic interpolated propagation scheme for solving the hyper-dimensional Vlasov-Poisson equation in phase space. *Computer Physics Communications*, 120(2):122–154, 1999.

- [42] J. Denavit. Numerical simulation of plasmas with periodic smoothing in phase space. *Journal of Computational Physics*, 9(1):75–98, 1972.
- [43] F. Filbet and E. Sonnendrücker. Numerical methods for the Vlasov equation. In *Numerical mathematics and advanced applications*, pages 459–468. Springer, 2003.
- [44] A.J. Klimas. A method for overcoming the velocity space filamentation problem in collisionless plasma model solutions. *Journal of computational physics*, 68(1):202–226, 1987.
- [45] AJ Klimas and WM Farrell. A splitting algorithm for Vlasov simulation with filamentation filtration. *Journal of computational physics*, 110(1):150–163, 1994.
- [46] Bengt Eliasson. Outflow boundary conditions for the Fourier transformed two-dimensional Vlasov equation. *Journal of Computational Physics*, 181(1):98–125, 2002.
- [47] Bengt Eliasson. Outflow boundary conditions for the Fourier transformed three-dimensional Vlasov–Maxwell system. *Journal of Computational Physics*, 225(2):1508–1532, 2007.
- [48] Frederick C Grant and Marc R Feix. Fourier-Hermite solutions of the Vlasov equations in the linearized limit. *Physics of Fluids (1958-1988)*, 10(4):696–702, 1967.
- [49] J.P. Holloway. Spectral velocity discretizations for the Vlasov-Maxwell equations. *Transport theory and statistical physics*, 25(1):1–32, 1996.
- [50] Glenn Joyce, Georg Knorr, and Homer K Meier. Numerical integration methods of the Vlasov equation. *Journal of Computational Physics*, 8(1):53–63, 1971.
- [51] Magdi Shoucri and Georg Knorr. Numerical integration of the Vlasov equation. *Journal of Computational Physics*, 14(1):84–92, 1974.
- [52] John V Shebalin. A spectral algorithm for solving the relativistic Vlasov–Maxwell equations. *Computer physics communications*, 156(1):86–94, 2003.
- [53] Oscar Buneman. Dissipation of currents in ionized media. *Physical Review*, 115(3):503, 1959.
- [54] John Dawson. One-dimensional plasma model. *Physics of Fluids (1958-1988)*, 5(4):445–459, 1962.
- [55] O Buneman and G Kooyers. Computer simulation of the electron mixing mechanism in ion propulsion. *AIAA Journal*, 1(11):2525–2528, 1963.
- [56] O Buneman and DA Dunn. Computer experiments in plasma physics. 1965.
- [57] Charles K Birdsall and A Bruce Langdon. *Plasma physics via computer simulation*. CRC Press, 2004.
- [58] Roger W Hockney and James W Eastwood. *Computer simulation using particles*. CRC Press, 1988.
- [59] H Neunzert and J Wick. The convergence of simulation methods in plasma physics. *Mathematical methods of plasmaphysics (Oberwolfach, 1979)*, 20:271–286, 1980.
- [60] G-H Cottet and P-A Raviart. Particle methods for the one-dimensional Vlasov-Poisson equations. *SIAM journal on numerical analysis*, 21(1):52–76, 1984.

- [61] HD Victory, Jr and Edward J Allen. The convergence theory of particle-in-cell methods for multidimensional Vlasov-Poisson systems. *SIAM journal on numerical analysis*, 28(5):1207–1241, 1991.
- [62] Stephen Wollman. On the approximation of the Vlasov–Poisson system by particle methods. *SIAM journal on numerical analysis*, 37(4):1369–1398, 2000.
- [63] Mickaël Melzani, Christophe Winisdoerffer, Rolf Walder, Doris Folini, Jean M Favre, Stefan Krastanov, and Peter Messmer. Apar-t: code, validation, and physical interpretation of particle-in-cell results. *Astronomy & Astrophysics*, 558:A133, 2013.
- [64] P Londrillo, A Sgattoni, and F Rossi. The aladyn code technical report. 2011.
- [65] Allen Taflove and Susan C Hagness. *Computational electrodynamics*. Artech house publishers, 2000.
- [66] Jian-Ming Jin. *The finite element method in electromagnetics*. John Wiley & Sons, 2014.
- [67] Ursula Van Rielen. *Numerical methods in computational electrodynamics: linear systems in practical applications*, volume 12. Springer Science & Business Media, 2012.
- [68] Kane S Yee et al. Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media. *IEEE Trans. Antennas Propag*, 14(3):302–307, 1966.
- [69] Rachel Nuter, Mickael Grech, Pedro Gonzalez de Alaiza Martinez, Guy Bonnaud, and Emmanuel d’Humières. Maxwell solvers for the simulations of the laser-matter interaction. *The European Physical Journal D*, 68(6):1–9, 2014.
- [70] T. Zh. Esirkepov. Exact charge conservation scheme for particle-in-cell simulation with an arbitrary form-factor. *Computer Physics Communications*, 135(2):144–153, 2001.
- [71] John Villasenor and Oscar Buneman. Rigorous charge conservation for local electromagnetic field solvers. *Computer Physics Communications*, 69(2-3):306–316, 1992.
- [72] J Candy and W Rozmus. A symplectic integration algorithm for separable hamiltonian functions. *Journal of Computational Physics*, 92(1):230–256, 1991.
- [73] Richard E Denton and M Kotschenreuther. δf algorithm. *Journal of Computational Physics*, 119(2):283–294, 1995.
- [74] William B Bateson and Dennis W Hewett. Grid and particle hydrodynamics:: Beyond hydrodynamics via fluid element particle-in-cell. *Journal of Computational Physics*, 144(2):358–378, 1998.
- [75] TD Arber, K Bennett, CS Brady, A Lawrence-Douglas, MG Ramsay, NJ Sircombe, P Gillies, RG Evans, H Schmitz, AR Bell, et al. Contemporary particle-in-cell approach to laser-plasma modelling. *Plasma Physics and Controlled Fusion*, 57(11):113001, 2015.
- [76] B Svedung Wettervik, TC DuBois, and T Fülöp. Vlasov modelling of laser-driven collisionless shock acceleration of protons. *Physics of Plasmas (1994-present)*, 23(5):053103, 2016.
- [77] Anna Grassi, Luca Fedeli, Andrea Sgattoni, and Andrea Macchi. Vlasov simulation of laser-driven shock acceleration and ion turbulence. *Plasma Physics and Controlled Fusion*, 58(3):034021, 2016.

- [78] S. Sgattoni, L. Fedeli, S. Sinigardi, and A. Marocchino. `piccante`: a spicy massively parallel fully-relativistic electromagnetic 3d particle-in-cell code. <http://aladyn.github.io/piccante/>, 2015.
- [79] T.A. Witten and L.M. Sander. Diffusion-limited aggregation, a kinetic critical phenomenon. *Phys. Rev. Lett.*, 47:1400–1403, Nov 1981.
- [80] Tatsufumi Nakamura, Motonobu Tampo, Rryosuke Kodama, Sergei V Bulanov, and Masaki Kando. Interaction of high contrast laser pulse with foam-attached target. *Physics of Plasmas (1994-present)*, 17(11):113107, 2010.
- [81] L. Fedeli, A. Sgattoni, G. Cantono, D. Garzella, F. Réau, I. Prencipe, M. Passoni, M. Raynaud, M. Květoň, J. Proska, A. Macchi, and T. Ceccotti. Electron acceleration by relativistic surface plasmons in laser-grating interaction. *Phys. Rev. Lett.*, 116:015001, Jan 2016.
- [82] A. Sgattoni, L. Fedeli, G. Cantono, T. Ceccotti, and A. Macchi. High field plasmonics and laser-plasma acceleration in solid targets. *Plasma Physics and Controlled Fusion*, 58(1):014004, 2016.
- [83] A. Sgattoni, S. Sinigardi, L. Fedeli, F. Pegoraro, and A. Macchi. Laser-driven rayleigh-taylor instability: Plasmonic effects and three-dimensional structures. *Phys. Rev. E*, 91:013106, Jan 2015.
- [84] A. Sgattoni, S. Sinigardi, and A. Macchi. High energy gain in three-dimensional simulations of light sail acceleration. *Applied Physics Letters*, 105(8), 2014.
- [85] A. Sgattoni, L. Fedeli, S. Sinigardi, A. Marocchino, A. Macchi, V. Weinberg, and A. Karmakar. Optimising PICCANTE - an open source particle-in-cell code for advanced simulations on tier-0 systems. *CoRR*, abs/1503.02464, 2015.
- [86] EL Clark, K Krushelnick, M Zepf, FN Beg, M Tatarakis, A Machacek, MIIK Santala, I Watts, PA Norreys, and AE Dangor. Energetic heavy-ion and proton generation from ultraintense laser-plasma interactions with solids. *Physical Review Letters*, 85(8):1654, 2000.
- [87] T. Lin, K. Flippo, M. Rever, A. Maksimchuk, and Donald P. Umstadter. Mechanism and control of high-intensity-laser-driven proton acceleration. 2004.
- [88] J. Metzkes, T. Kluge, K. Zeil, M. Bussmann, S. D. Kraft, T. E. Cowan, and U. Schramm. Experimental observation of transverse modulations in laser-driven proton beams. *New Journal of Physics*, 16(2):023008, 2014.
- [89] The Message Passing Interface (MPI) standard. <http://www.mpi-forum.org/> and <http://www.mpi-forum.org/>.
- [90] OpenMP: API specification for parallel programming. <http://openmp.org/>.
- [91] Viktor K Decyk and Tajendra V Singh. Adaptable particle-in-cell algorithms for graphical processing units. *Computer Physics Communications*, 182(3):641–648, 2011.
- [92] Heiko Burau, Renée Widera, Wolfgang Honig, Guido Juckeland, Alexander Debus, Thomas Kluge, Ulrich Schramm, Tomas E Cowan, Roland Sauerbrey, and Michael Bussmann. Picon-gpu: A fully relativistic particle-in-cell code for a gpu cluster. *IEEE Transactions on Plasma Science*, 38(10):2831–2839, 2010.

- [93] Sergey Bastrakov, Iosif Meyerov, Igor Surmin, Evgeny Efimenko, Arkady Gonoskov, Alexander Malyshev, and Mikhail Shiryaev. Particle-in-cell plasma simulation on cpus, gpus and xeon phi coprocessors. In *ISC*, pages 513–514. Springer, 2014.
- [94] Michael A Purvis, Vyacheslav N Shlyaptsev, Reed Hollinger, Clayton Bargsten, Alexander Pukhov, Amy Prieto, Yong Wang, Bradley M Luther, Liang Yin, Shoujun Wang, et al. Relativistic plasma nanophotonics for ultrahigh energy density physics. *Nature Photonics*, 7(10):796–800, 2013.
- [95] Takayuki Umeda, Yoshiharu Omura, T Tominaga, and Hiroshi Matsumoto. A new charge conservation method in electromagnetic particle-in-cell simulations. *Computer Physics Communications*, 156(1):73–85, 2003.
- [96] Smilei. <http://smileipic.github.io/Smilei/>.