



HAL
open science

Massively parallel algorithms for realistic PIC simulations of ultra high intensity laser-plasma interaction, application to attosecond pulses separation of Doppler harmonics

Haithem Kallala

► **To cite this version:**

Haithem Kallala. Massively parallel algorithms for realistic PIC simulations of ultra high intensity laser-plasma interaction, application to attosecond pulses separation of Doppler harmonics. Computational Physics [physics.comp-ph]. Université Paris-Saclay, 2020. English. NNT : 2020UPASS052 . tel-02910065

HAL Id: tel-02910065

<https://theses.hal.science/tel-02910065>

Submitted on 31 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Massively parallel algorithms for realistic PIC simulations of ultra high intensity laser-plasma interaction, application to attosecond pulses separation of Doppler harmonics

Thèse de doctorat de l'Université Paris-Saclay

École doctorale n° 572, Ondes et Matière, EDOM
Spécialité de doctorat: Physique des plasmas
Unité de recherche: Université Paris-Saclay, CEA, CNRS, LIDYL, 91191,
Gif-sur-Yvette,
Réfèrent: : Faculté des sciences

Thèse présentée et soutenue à Saclay, le 21/02/2020, par

Haithem Kallala

Composition du jury:

Gilles Maynard Directeur de recherche CNRS, UMR 8578	Président
Emmanuel D'Humières Professeur des universités, Université de Bordeaux	Rapporteur & Examineur
Eric Sonnendrücker Professeur, Université technique de Munich	Rapporteur & Examineur
Mickael Grech Chargé de recherche, LULI, UMR 7605	Examineur
Caterina Riconda Professeure des universités, UPMC	Examinatrice
Pascal Monot Chercheur, LIDYL	Directeur de thèse
Henri Vincenti Chercheur, LIDYL	Invité
Julien Derouillat Chercheur, MDLS	Invité

Remerciements

Ce travail de thèse a été réalisé à la Maison De La Simulation (MDLS) et au sein de l'équipe de physique des ultra-hautes intensités (PHI) du Laboratoire Interactions, Dynamiques et Lasers (LIDYL). Je remercie Edouart Audit ainsi que Philippe Martin pour m'avoir accueilli dans leur laboratoires respectifs pendant plus de trois années.

Mes remerciements vont à Pascal Monot, mon directeur de thèse, qui, tout en me conférant une grande liberté de travail, a toujours su bien veillé au bon avancement de mes travaux. Je remercie également Fabien Quéré, responsable du groupe PHI, pour ses grandes qualités scientifiques et humaines, et pour l'excellente ambiance qu'il veille à faire régner au sein du groupe.

J'aimerais aussi exprimer ma gratitude à Gilles Maynard pour l'intérêt qu'il a témoigné pour ce travail et pour avoir accepté de présider le jury de thèse. Mes remerciements vont aussi à Eric Sonnendrücker et Emmanuël D'Humières pour avoir accepté d'être les rapporteurs de ce travail, ainsi qu'à Mickael Grech et Caterina Riconda pour leur participation au jury.

Je remercie chaleureusement mon encadrant au LIDYL, Henri Vincenti. Durant ces trois années, j'ai grandement appris de son expertise multi-disciplinaire très large qui va de la physique des plasmas, à la programmation scientifique, en passant par l'analyse numérique, et qu'il n'hésite pas à partager avec tous ses étudiants. Ses intuitions physiques et numériques et sa pédagogie ont largement guidé mes travaux de thèse. Je te remercie également pour ta grande patience pendant la rédaction du manuscrit et pour toutes les relectures très minutieuses qui l'ont grandement amélioré. Ces trois années à travailler ensemble m'ont permis d'apprécier tes grandes qualités humaines à chaque fois qu'on se réunit autour d'un déjeuner ou lorsqu'on prend le bus de 19:56 pour rentrer à Massy.

Je remercie aussi mon encadrant à la MDLS, Julien Derouillat, pour sa confiance et sa bienveillance. J'ai tant appris de toi sur le plan humain et scientifique. Merci pour tes encouragements pendant les moments les plus délicats et pour tous les conseils et tes remarques qui ont rajouté de la cohérence et de la pertinence à ce travail. Mes remerciements vont aussi à nos collaborateurs du LBNL: Jean-Luc Vay, Rémi Lehe, Maxence Thévenet.

J'aimerais remercier par ailleurs toutes les personnes avec qui j'ai pu travailler ou partager des moments agréables pendant cette thèse. Je remercie vivement l'ensemble du personnel de la Maison de La Simulation: Julien Bigot, Nathalie, Maxime, Ksander, Arnaud, Valérie, Pascal, Yacine, Anders, Marion, Axelle, Layla.

Je remercie par ailleurs tous les membres du LIDYL. Une pensée particulière va naturellement à l'équipe PHI avec qui j'ai partagé une expérience professionnelle et humaine extraordinaire. Je remercie tous mes collègues de bureau qui ont dû supporté la nuisance sonore que fait mon ordinateur lorsque j'y fais tourner des simulations. Merci à Luca Fedeli, Pierre Forestier-Colleoni, Antonin Sainte-Marie, Cedric Bome et Murad Abuzarli. Je remercie aussi les autres thésards de l'équipe, Guillaume Bouchard, Antoine Jeandet et Guillaume Blaclard pour tous les débats animés qu'on tenait régulièrement à la cantine ou autour de la machine à café, et qui nous ont permis de nouer une amitié sincère au-delà du travail.

Je n'aurai pu aller au bout de ce travail sans le soutien de bon nombre de personnes en dehors du CEA. Je remercie en premier lieu mes parents pour leurs encouragements et leurs soutiens inconditionnel pendant la thèse et pour les valeurs essentielles de dévouement et de persévérance qu'ils m'ont inculqué depuis mon enfance. Je remercie aussi mes quatres frères, Amine, Khalil, Hassan et Youssef pour leur encouragements ainsi que ma tante Najet et mon oncle Salem qui ont fait le déplacement à Paris pour la soutenance. Je remercie aussi mes cousins Parisiens Mohamed et Ines pour les meetings familiaux réguliers qui me permettaient de décompresser. Je remercie finalement ma copine, Imen, qui connaît probablement le mieux la face cachée de cette thèse, pour m'avoir supporté durant les moments les plus difficiles de ce travail de longue haleine.

Table of contents

I	Résumé en français	1
II	Introduction	4
II.1	Context	4
II.2	Objectives and Outline	6
A	State of the art of high order harmonic generation and Particle-In-Cell codes.	9
III	Physics of high order harmonics generation on plasma mirrors	10
III.1	The plasma state	11
III.1.1	Langmuir frequency	11
III.1.2	Electromagnetic wave propagation in a plasma	12
III.2	Gaussian laser formalism	14
III.3	High harmonic generation on plasma mirrors	16
III.3.1	Plasma mirror Formation	16
III.3.2	Harmonic generation mechanisms	18
III.3.3	Attosecond pulses emission	18
III.4	Models for HHG on plasma mirrors	20
III.4.1	Bourdier boosted frame	20
III.4.2	Wave equation with sources in Bourdier frame	22
III.4.3	Transverse currents expression	23
III.4.4	Coherent wake emission regime	24
III.4.5	Harmonic generation in the relativistic regime	25
III.4.6	Doppler effect	25
III.4.7	The Relativistic Oscillating Mirror (ROM) model	26
III.4.8	Plasma denting model and spatial properties of Doppler harmonics	28
III.4.9	Plasma mirror excursion in three dimensions:	30
III.4.10	Spatial properties of high order harmonics	31
III.5	The attosecond lighthouse effect	34
III.5.1	Characterization of ultra-short laser pulses	35
III.5.2	First order spatio-temporal couplings in ultra short lasers	36
III.5.3	Different representations of spatio-temporal couplings	37

III.5.4	Practical realization of the attosecond lighthouse effect	39
III.5.5	Inducing pulse-front tilt	41
III.5.6	From PFT to WFR	43
III.5.7	Separation condition of attosecond light pulses	45
IV	Basics of the PIC algorithm	47
IV.1	Numerical modelling in plasma physics	48
IV.2	Maxwell-Vlasov equations	48
IV.3	The Particle-In-Cell (PIC) method	49
IV.3.1	WARP + PICSAR PIC framework	50
IV.3.2	Discretization of PIC quantities	51
IV.3.3	Current/charge deposition	52
IV.3.4	Fields interpolation	53
IV.3.5	Particles pusher	54
IV.3.6	Maxwell Solver	54
IV.4	FFT-based Pseudo-spectral Maxwell solvers	60
IV.4.1	Pseudo-Spectral Analytical Time Domain	60
IV.4.2	Parallelization of the pseudo-spectral methods	63
IV.4.3	Truncation error and ultra-high order solvers	63
IV.5	Absorbing boundary conditions with PSATD	66
B	Development of numerical tools for accurate ultra high intensity physics simulations	69
V	A generalized massively parallel scheme for FFT-based pseudo-spectral Maxwell solvers	70
V.1	Introduction	71
V.2	Brief description of HPC systems	72
V.3	Scalability limits of existing parallelization techniques for pseudo-spectral solvers	73
V.3.1	Scalability limits of the global method	73
V.3.2	Scalability limits of the local method	77
V.4	A generalized massively parallel FFT-based Maxwell solver	79
V.4.1	Principle of the technique	79
V.4.2	Advantages of the new parallelization strategy	80
V.4.3	Coupling of the hybrid solver with the full PIC algorithm: a dual grid decomposition for efficient load balancing	86
V.4.4	Benchmarks of the hybrid solver on different architectures	89
V.5	Conclusion and prospects	93

VI Boosted frame as a visualization tool for UHI physics	95
VI.1 Context	96
VI.2 Need for an efficient visualization tool in UHI physics	96
VI.2.1 Limitations of HHG emission visualization in the laboratory frame	96
VI.3 Boosted frame visualization for PIC-UHI simulations	99
VI.3.1 Lorentz transformation in PIC codes	99
VI.4 High performance on the fly implementation of Lorentz transformation	101
C Physics study with the new tools	107
VII Generating isolated attosecond pulses in the relativistic regime	108
VII.1 Introduction	109
VII.2 Limitations of the attosecond lighthouse effect for Doppler harmonics	112
VII.2.1 Physical setup	112
VII.2.2 Numerical parameters	112
VII.2.3 Post-processings	114
VII.2.4 Results	115
VII.3 Effect of the plasma density scale length on the harmonic beam divergence	117
VII.4 PPM curvature compensation via laser wavefront curvature	120
VII.4.1 General principle	120
VII.4.2 Divergence of high order harmonics generated by a spatially chirped and defocused laser pulse	121
VII.4.3 Numerical validation	128
VII.4.4 Evolution of the optimal defocusing length with the gradient scale length	130
VII.4.5 Effect of the laser incidence angle	133
VII.4.6 Effect of the pulse-front tilt on the angular separation of attosecond light pulses	135
VII.4.7 Attosecond light pulses separation in 3D	141
VII.5 Laser pulse shaping	148
VII.5.1 General principle	148
VII.5.2 Flattening a laser pulse profile in the presence of spatial chirp	149
VII.5.3 PIC simulations of ROM harmonics generation with a shaped laser pulse	152
VII.6 Conclusion	155
VIII Conclusions and perspectives	156
A Derivation of laser parameters out of focus in the presence of wavefront rotation	158
B Relationship between spatio-spectral beam sizes	163
Bibliography	164

Chapter I

Résumé en français

Cette thèse porte sur le développement d'algorithmes massivement parallèle pour la physique des ultra-hautes intensités, et sur l'étude de l'interaction laser-miroir plasma, dans le but de produire des impulsions attosecondes isolées en régime relativiste.

La complexité des mécanismes physiques mis en jeu lors de l'interaction laser-plasma à ultra-haute intensité nécessite de recourir à des simulations PIC particulièrement lourdes. Au cœur de ces codes de calcul, la résolution numérique des équations de Maxwell à chaque pas de temps constitue une étape très importante et conditionne fortement la qualité et la précision des résultats numériques obtenus. Dans beaucoup de cas de figure d'intérêt en physique des hautes intensités, les algorithmes numériques standards, de type différence finie, peuvent induire diverses erreurs liées à la dispersion et à l'anisotropie numérique inhérentes à ce type d'approche. Ce qui peut compromettre la robustesse de la modélisation.

Des solveurs de Maxwell pseudo-spectraux d'ordre élevé présentent de nombreux avantages en termes de précision numérique par rapport aux approches standard. Néanmoins, ces solveurs ont un coût élevé en termes de temps de calcul et de ressources de calcul nécessaires. En effet, les techniques de parallélisation existant pour ces solveurs sont peu performantes au-delà de quelques milliers de cœurs, ou induisent un important usage mémoire, ce qui limite leur scalabilité à large échelle. Dans cette thèse, nous avons développé une toute nouvelle approche de parallélisation hiérarchique qui combine les avantages des méthodes existantes. Cette méthode a été testée à très large échelle et montre un scaling significativement meilleur que les précédentes techniques, tout en garantissant une occupation mémoire réduite.

En complément à ce travail numérique, j'ai développé un nouvel outil de diagnostic numérique massivement parallèle basé sur des changements référentiels Lorentziens. Le changement de référentiel relativiste présente de multiples intérêts en physique UHI. En effet, il s'agit d'un outil de modélisation théorique important qui permet de décrire plus naturellement les processus physiques en jeu lors de l'interaction laser-miroir plasma en incidence oblique, en découplant les couplages spatio-temporels sous-jacents à l'incidence oblique. C'est aussi un outil algorithmique important qui permet de réduire les temps de simulation de plusieurs ordres de grandeur dans le cadre des études portant sur

l'accélération par sillage laser dans les plasmas sous denses.

En capitalisant sur ce travail numérique et algorithmique, nous avons réalisé une étude numérique et théorique approfondie dans le cadre de la génération d'harmoniques d'ordre élevé sur cible solide. Lorsqu'une impulsion laser ultra-intense ($I > 10^{16} \text{Wcm}^{-2}$) et ultra-courtes (de quelques dizaines de femtosecondes) est focalisée sur une cible solide, elle génère un plasma sur-dense, appelé miroir plasma, qui réfléchit non linéairement le laser incident. La réflexion de l'impulsion laser est accompagnée par l'émission cohérente d'harmoniques d'ordre élevée, sous forme d'impulsions X-UV attosecondes (10^{-18}s). Pour des intensités laser relativiste ($I > 10^{19} \text{Wcm}^{-2}$), la surface du plasma est incurvée sous l'effet de la pression de radiation du laser. De ce fait, les harmoniques rayonnées par la surface du plasma sont focalisées. Dans cette thèse, j'ai étudié la possibilité de produire des impulsions attosecondes isolées en régime relativiste sur miroir plasma, grâce au mécanisme de phare attoseconde. Celui-ci consiste à introduire une rotation des fronts d'onde du laser incident de façon à séparer angulairement les différentes impulsions attosecondes produites à chaque cycle optique. En régime relativiste, la courbure du miroir plasma augmente considérablement la divergence du faisceau harmonique, ce qui rend le mécanisme phare attoseconde inefficace. Pour y remédier, j'ai développé deux techniques de réduction de divergence harmonique afin de mitiger l'effet de focalisation induit par la courbure du miroir plasma et permettre de générer des impulsions attosecondes isolées à partir d'harmoniques Doppler. Ces deux techniques sont basés sur la mise en forme en amplitude et en phase du faisceau laser. Par ailleurs, j'ai développé un modèle théorique pour déterminer les régimes optimaux d'interaction afin de maximiser la séparation angulaire des impulsions attosecondes. Ce modèle a été validé par des simulations numériques PIC en géométries 2D et 3D et sur une large gamme de paramètres laser et plasma. Finalement, on montre qu'en ajustant des paramètres laser et plasma réalistes, il est possible de séparer efficacement les impulsions attosecondes en régime relativiste.

Ce manuscrit est structuré comme suit:

La première partie donne une vue d'ensemble des mécanismes d'interaction laser ultra-intense et ultra-court avec de la matière et des outils numériques utilisés durant cette thèse. Le premier chapitre passe en revue les différents mécanismes de génération d'harmoniques sur plasma sur-dense, et introduit la théorie des couplages spatio-temporels des lasers ultra-courts. Nous détaillons ensuite l'état de l'art de la technique phare attosecond développée il y a quelques années en vue de produire des impulsions attosecondes isolées spatialement. Cette technique se base sur l'introduction de rotation de fronts d'onde du laser incident. Ainsi, les impulsions attosecondes successives générées à la surface du miroir plasma à chaque cycle optique se propagent selon des directions légèrement différentes. Si ses impulsions sont suffisamment bien collimatées, elles seront séparées angulairement loin de la cible. Le deuxième chapitre introduit la méthode Particle-In-Cell, utilisée pour réaliser des simulations d'interaction laser-plasma. Nous discutons des limites des solveurs standards pour les équations de Maxwell et introduisons l'algorithme de résolution pseudo-spectral (PSATD) ainsi que la méthode de parallélisation la plus efficace proposée dans la littérature. Nous présentons par ailleurs un formalisme

pour rendre cet algorithme compatible avec des conditions de bord absorbante.

La deuxième partie est consacrée aux développements algorithmiques et numériques réalisés au cours de cette thèse. Dans le premier chapitre, nous discutons des limites des méthodes de parallélisation existant pour les solveurs de Maxwell pseudo-spectraux et introduisons notre nouvelle approche, dit hybride. Un modèle de scalabilité à très large échelle, prenant en compte les performances du réseau de la machine de calcul, est dérivé. Nous montrons via une série de hêtre-marks réalisés à très large échelle que cette approche apporte un gain important en matière de temps de calcul et d'occupation mémoire.

Le deuxième chapitre de cette partie présente la stratégie d'implémentation massivement parallèle de l'outil de diagnostic basé sur le changement de référentiel Lorentzien. On montre, via des simulations, que les propriétés spatiales des harmoniques issues du miroir plasma dans différents régimes sont naturellement mieux décrites dans ce référentiel.

La troisième partie est consacrée à l'étude de l'effet phare attoseconde en régime relativiste. Nous discutons tout d'abord les limites de cette technique en régime relativiste. Nous présentons ensuite deux techniques de réduction de divergences harmoniques, basé sur un contrôle réaliste en phase et en intensité du laser incident pour permettre une séparation angulaire efficace des impulsions attosecondes loin de la cible. Ces deux techniques sont étudiés numériquement à travers des simulations massivement parallèles en géométrie 2D et 3 d. Les résultats des simulations montrent que des impulsions attosecondes uniques peuvent être obtenues avec des paramètres laser et plasma réalistes en régime Peta Watt.

Finalement, des perspectives générales sur la portée des travaux réalisés durant cette thèse sont donnés en conclusion.

Chapter II

Introduction

II.1 Context

The emergence and subsequent formidable development of light amplification technologies during the second half of the twentieth century made coherent light sources available for a wide panel of scientific and societal uses. With the invention of chirped pulse amplification techniques (CPA) [1, 2], modern laser systems are capable of delivering extremely energetic ($10^{-3} - 1000$ J) light pulses, within ultra-short time scales (4 femtoseconds - 1 picoseconds). The use of these systems for scientific purposes has opened up a new and vastly unexplored research branch commonly called ultra-high intensity (UHI) physics. When focused on matter, the electromagnetic field intensities produced by these lasers, reaching up to $10^{22} W.cm^{-2}$, turn matter into out of equilibrium plasmas, predominantly governed by non-linear and relativistic effects [3]. UHI physics concentrates on the study of laser and matter interactions under extreme irradiation conditions and within ultra-short time scales for fundamental and applicative perspectives.

In recent years, a lot of the research effort is driven towards the perspective of developing future high energy particles and/or photons sources based on laser-matter interaction schemes. In this context, the so-called laser wakefield acceleration technique [4] for example, is aiming at developing laser-based electron accelerators. An ultra-intense laser pulse is used to excite large amplitude plasma waves (up to $100 GV.m^{-1}$) inside an underdense plasma resulting from the ionization of a gaseous medium. These plasma waves act as accelerating structures that can boost electrons to multi-Gev energy ranges within few centimeters length scales [5, 6, 7, 8]. This acceleration setup may have a wide variety of applications, including compact particle colliders and high energy x-ray sources [9].

Another hot topic in the context of UHI physics is the study of solid target, ultra-intense laser interaction. In this case, the laser instantly turns the solid target surface into a fully ionized overdense plasma medium, commonly called a plasma mirror (PM). Plasma mirrors reflect off the incident laser pulse. This reflection implies highly non-linear mechanisms that may result in the coherent emission of high order laser harmonics in the form of sub-femtosecond (ie. attosecond) X-UV light pulses [10, 11] [12, 13, 14, 15, 16]. Laser interaction with solid media can also be leveraged to produce multi-Mev sources of electrons [17] and ions [18].

The use of plasma mirrors as a source of bright X-UV radiation may constitute a potential alternative to standard attosecond light sources for pump-probe experiments and ultra-fast spectroscopy. Indeed, in contrast with standard attosecond pulses generation techniques in gaseous media [19, 20, 21, 22], which are intrinsically limited to medium-low laser intensities, high order harmonic generation (HHG) on plasma mirrors can exploit the full capabilities of ultra-high intense lasers to achieve unmatched harmonic brightness.

However, many difficulties still need to be addressed before achieving this goal. Those limitations include the currently low repetition rates of TeraWatt and PetaWatt class lasers (0.01-10 Hz) and the difficulty of achieving a strong control degree in HHG mechanisms on PMs in order to produce unique or isolated attosecond light pulses, naturally more suited for applications. The latter milestone is the subject of an in-depth theoretical and numerical study in this thesis, where we identify optimal interaction regimes to produce isolated attosecond pulses from plasma mirror harmonics.

With the advent of scientific computing, ultra-high intensity physics heavily relies on numerical simulations to corroborate theoretical predictions, complete experimental studies or investigate and imagine unexplored interaction regimes.

UHI physics essentially exploits Particle-In-Cell (PIC) codes to conduct ab initio simulations of laser-plasma interactions. The PIC algorithm is a particle-mesh kinetic approach to simulate plasma dynamics by self-consistently modelling the evolution of electromagnetic fields and free particles dynamics.

Laser-plasma interaction scenarios often involve complex and multi-scale mechanisms, implying ultra-energetic particles and extremely short electromagnetic wavelengths. Quantitatively reproducing all relevant mechanisms taking place in laser-plasma problems with PIC algorithms usually mandates very fine mesh samplings due to the presence of rapidly varying electromagnetic fields and a huge number of numerical particles to finely describe statistical effects at play. For instance, a realistic simulation of HHG on plasma mirrors for example in a 3D geometry can require hundreds of billions of particles and an equivalent amount of mesh cells. Therefore, numerical simulations in UHI physics pose two mutually interdependent numerical difficulties.

First, such simulations mandate huge computational resources, both in terms of memory occupancy and computing power, and can only be performed on massively parallel super-computers. After nearly half a century of extensive development and optimization, the most advanced computers available today can achieve hundreds of Petaflops ¹ (10^{15} floating-point operations per second), while the Exaflop barrier is expected to be reached in the next few years. These incredible performances were made possible by implementing increasingly complex hardware architectures. Modern parallel computing paradigms favor hierarchical parallelisms to achieve cost-effective computations. State-of-art

¹<https://www.top500.org/system/179397>

massively parallel platforms are based on a distributed memory setup, where multiple compute nodes are connected with a high bandwidth network. Individual compute nodes enable additional shared memory multiple-instructions-multiple-data parallel capabilities by leveraging many-core or multi-core architectures, and where each computing core is made of multiple vector units to execute shared instructions on multiple data items (SIMD). Compute node may be complemented by accelerating coprocessing devices (such as a GP-GPU or FPGA ...) in heterogeneous systems. Efficient PIC codes should be well adapted to these computing architectures. In this regard, when developing a PIC code, great care should be given to scalability performances and parallelization strategies in use in order to access the full capacities of computing resources.

The second milestone is related to the accuracy of the numerical methods employed within the PIC algorithm. Those methods are expected to robustly solve the dynamics of such multi-scale and complex systems as plasma irradiated by lasers. Indeed, identifying and inhibiting spurious numerical artifacts is of critical importance for PIC simulations. Various purely numerical instabilities that severely hinder the accuracy of PIC simulations for UHI laser-plasma simulations result from the unphysical numerical dispersion relations associated with standard discretized Maxwell's equations solvers. Such instabilities include the well known numerical Cherenkov effect [23] in the context of laser wakefield acceleration simulations or the recently identified unphysical refraction of high order harmonics emitted by plasma mirrors [24].

Simultaneously meeting the scalability and the accuracy challenges mandate the development of sophisticated numerical methods along with optimized parallelization strategies. In the context of PIC simulations, this often implies a trade-off between numerical accuracy and computational performance. Indeed, scaling standard Finite Difference Time Domain (FDTD) Maxwell's equations solvers up to millions of cores is relatively straightforward as it would only require a regular Cartesian domain decomposition to evenly split the workload between different compute nodes and limited inter-node data communications. However, PIC simulations employing finite difference Maxwell solvers remain prone to various numerical instabilities due to the important numerical dispersion of FDTD methods.

More robust numerical approaches encompass the use of high order stencils or FFT-based pseudo-spectral solvers to mitigate numerical errors. However, those methods mandate important inter-node communications, which often result in significant network saturation, limiting their scalability to moderate scales only. Therefore, designing highly scalable parallelization strategies for this type of solvers is of significant importance to enable their use in the context of massively parallel PIC-UHI simulations.

II.2 Objectives and Outline

This thesis lies between the fields of computational physics and High Performance Computing. At the start of my PhD, an innovative FFT-based, massively parallel, Maxwell solver algorithm had recently been proposed and benchmarked at Lawrence Berkeley National Laboratory [25, 24]. This

algorithm may, however, suffer from an important memory footprint and a scalability loss at very large scales. I first was involved in porting this algorithm into the SMILEI² PIC code developed at Maison De La Simulation at CEA. Then, I implemented a novel parallelization strategy that results in a substantial speed-up, while decreasing the memory footprint by one order of magnitude. This novel parallelization approach was implemented in the PICSAR³ library, and can be used with WARP⁴ framework. I also actively contributed to various optimizations and implementations in the WARP+PICSAR framework. In this context, I developed a novel massively parallel numerical diagnostic tool for PIC simulation, which relies on Lorentz transformation to better describe the physics at play in the context of UHI simulations.

By leveraging those novel high scalable numerical tools to perform large scales and massively parallel PIC simulations, I conducted an extensive numerical and theoretical study in the context of high order harmonic generation on plasma mirrors. In particular, I was interested in the generation of isolated attosecond light pulses from plasma mirrors in the relativistic regime. In this regard, an innovative scheme to angularly separate attosecond light pulses has been proposed at CEA [26] a few years ago. However, this approach cannot be directly transposed to the relativistic regime, because the optimal setup that maximizes the harmonic generation efficiency increases the harmonic beam divergence. My work involved identifying and characterizing optimal interaction conditions to enable the angular isolation of attosecond light pulses. For this matter, I developed an analytical toy model to predict the angular separation of attosecond light pulses based on the laser and plasma parameters, and in the presence of spatio-temporal couplings. This model was validated with an extensive numerical parametric study in 2D and 3D geometries, over a broad range of laser and plasma parameters. In the end, we designed two efficient all-optical schemes for generating isolated attosecond pulses, both involving specific control of the driving laser phase.

The second chapter presents some basic notions of laser-plasma interaction physics and details the essential aspects of high order harmonic generation on plasma mirrors, with a focus of Doppler harmonics generation.

The third chapter introduces the Particle-In-Cell codes that I used during my PhD. It also introduces the FFT-based Maxwell solver algorithm that is the heart of the development made in the third chapter.

The fourth chapter discusses the limitations of this Maxwell solver in terms of scalability and introduces the novel parallelization strategy developed during this thesis.

The fifth chapter presents the new massively parallel numerical diagnostics for PIC UHI simulations, also implemented in the WARP+PICSAR code. This versatile tool enables us to reconstruct

²<http://www.maisondelasimulation.fr/smilei/>

³<https://picsar.net/>

⁴<https://blast.lbl.gov/blast-codes-warp/>

simulations observables in a Lorentz boosted frame where data interpretation can be more easily carried out. This implementation has been benchmarked on different scenarios of high order harmonic generation on plasma mirrors. It shows that the physical pictures underlying the harmonic emission on plasma mirrors are better captured with this tool, compared to standard laboratory frame diagnostics.

The sixth chapter is dedicated to the generation of isolated attosecond light pulses in the relativistic regime. First, I highlight the limitations of the attosecond lighthouse techniques in this regime and identify the physical obstacles that hinder its applicability. I then conduct a numerical and theoretical study to identify optimal setups where attosecond pulses separation can be achieved. I propose two efficient schemes to generate angularly separated attosecond light pulses in the relativistic regime. Those techniques are extensively validated with PIC simulations in 2D and 3D geometries.

Part A

State of the art of high order harmonic generation and Particle-In-Cell codes.

Chapter III

Physics of high order harmonics generation on plasma mirrors

III.1 The plasma state	11
III.1.1 Langmuir frequency	11
III.1.2 Electromagnetic wave propagation in a plasma	12
III.2 Gaussian laser formalism	14
III.3 High harmonic generation on plasma mirrors	16
III.3.1 Plasma mirror Formation	16
III.3.2 Harmonic generation mechanisms	18
III.3.3 Attosecond pulses emission	18
III.4 Models for HHG on plasma mirrors	20
III.4.1 Bourdier boosted frame	20
III.4.2 Wave equation with sources in Bourdier frame	22
III.4.3 Transverse currents expression	23
III.4.4 Coherent wake emission regime	24
III.4.5 Harmonic generation in the relativistic regime	25
III.4.6 Doppler effect	25
III.4.7 The Relativistic Oscillating Mirror (ROM) model	26
III.4.8 Plasma denting model and spatial properties of Doppler harmonics	28
III.4.9 Plasma mirror excursion in three dimensions:	30
III.4.10 Spatial properties of high order harmonics	31
III.5 The attosecond lighthouse effect	34
III.5.1 Characterization of ultra-short laser pulses	35
III.5.2 First order spatio-temporal couplings in ultra short lasers	36
III.5.3 Different representations of spatio-temporal couplings	37
III.5.4 Practical realization of the attosecond lighthouse effect	39
III.5.5 Inducing pulse-front tilt	41
III.5.6 From PFT to WFR	43
III.5.7 Separation condition of attosecond light pulses	45

When an ultra-intense, ultra-short laser pulse, with a high temporal contrast is focused over a solid target, matter is quasi-instantly ionized, giving rise to a plasma mirror. Plasma mirrors act as non-linear optical systems for electromagnetic waves and simultaneously reflect and temporally distort incident laser light. These non-linearities give rise to remarkable X-UV light and relativistic electron sources that may find promising scientific, industrial and societal applications.

The first section of this chapter introduces some of the basic principles of plasma interaction with light. We show how the non-linear response of a plasma medium can give rise to a high order harmonic generation (HHG) in the form of attosecond pulses. The second part of this chapter is dedicated to spatio-temporal couplings (STC) in ultra-fast laser pulses and their potential applications in the context of high order harmonic generations on plasma mirrors.

III.1 The plasma state

The plasma state is usually referred to as the fourth fundamental state of matter, after the solid, liquid and gaseous phases. Unlike other matter states, which are constituted of atoms or molecules, plasmas are a collection of positively and negatively charged particles that exhibit a collective behavior through electromagnetic fields induced by free charged particles. Put together, two charged particles interact with one another through Coulomb potentials as isolated individual particles. However, in the presence of a large collection of particles, each one of them simultaneously interacts with many nearby particles. In this case, the Coulomb potential associated with a given particle is shielded out collectively by neighboring particles over a characteristic scale length called Debye length. Hence, each particle interacts individually with particles laying inside a Debye sphere. On larger spatial scales, the plasma dynamics is governed by collective charged particles behaviors.

III.1.1 Langmuir frequency

We consider an electrically neutral, homogeneous plasma initially at equilibrium that verifies the global neutrality for $t < 0$: $n_e = Zn_i$ with n_e and n_i are the electronic and ionic densities respectively, and Z is the ions charge number. At $t = 0$, the electrons located between $x = 0$ and $x = \xi$ are slightly displaced by a distance of ξ . An electrostatic force resulting from the charge separation emerges and tends to restore the initial equilibrium. The electrostatic field originating from this force can be computed by solving the Maxwell-Gauss equation and reads:

$$E(\xi(t), t) = \frac{n_e}{e\varepsilon_0} \xi(t) \quad (\text{III.1})$$

where e is the electron charge and ε_0 the vacuum permittivity.

The equation of motion of the displaced electrons is:

$$\frac{d^2\xi(t)}{dt^2} = -\frac{e}{m_e} E(\xi(t), t) = -\frac{n_e e^2}{m_e \varepsilon_0} \xi(t) \quad (\text{III.2})$$

where m_e the electronic mass. The solution of this equation is

$$\xi(t) = \xi(0)\cos(\omega_{pe}t) \quad (\text{III.3})$$

where $\omega_{pe} = \sqrt{\frac{n_e e^2}{m_e \epsilon_0}}$ is the so-called Langmuir frequency. This physical quantity is fundamental in plasma physics as it also plays an important role in plasma interaction with electromagnetic waves.

III.1.2 Electromagnetic wave propagation in a plasma

Let us consider monochromatic wave propagation in unmagnetized plasmas. The propagation is governed by Maxwell's equations:

$$\begin{aligned}\vec{\nabla} \cdot \vec{E} &= \frac{\rho}{\epsilon_0} \\ \vec{\nabla} \cdot \vec{B} &= 0 \\ \vec{\nabla} \wedge \vec{E} &= -\frac{d\vec{B}}{dt} \\ \vec{\nabla} \wedge \vec{B} &= \mu_0 \vec{J} + \frac{d\vec{E}}{c^2 dt}\end{aligned}\tag{III.4}$$

In Fourier space (k, ω) , Maxwell's equations write:

$$\begin{aligned}i\vec{k} \cdot \tilde{E} &= \frac{\tilde{\rho}}{\epsilon_0} \\ i\vec{k} \cdot \tilde{B} &= 0 \\ i\vec{k} \wedge \tilde{E} &= i\omega_0 \tilde{B} \\ i\vec{k} \wedge \tilde{B} &= \mu_0 \tilde{J} - \frac{i\omega_0 \tilde{E}}{c^2}\end{aligned}\tag{III.5}$$

As ions are much more massive than electrons, we assume immobile ions and only consider the currents induced by the electronic motions.

$$\tilde{J} = -n_e e \tilde{V}_e\tag{III.6}$$

Neglecting the magnetic force on electrons (assuming that $\|\vec{V}_e\| \ll c$) and assuming that the plasma is non collisional and cold, we have:

$$\tilde{V}_e = -\frac{e}{im_e \omega} \tilde{E}\tag{III.7}$$

Finally, the electric current is:

$$\tilde{J} = -\frac{\omega_{pe}^2}{\omega} \epsilon_0 \tilde{E}\tag{III.8}$$

Injecting the current expression into Maxwell's equations, we get:

$$(\vec{k} \cdot \tilde{E}) \vec{k} - \vec{k} \cdot \tilde{E} = \left(\frac{\omega_{pe}^2}{c^2} - \omega^2 \right) \tilde{E}\tag{III.9}$$

This is the propagation equation of an electromagnetic field in a medium with a refractive index:

$$N(\omega) = \sqrt{1 - \frac{\omega_{pe}^2}{\omega^2}}\tag{III.10}$$

Equation III.9 admits two kinds of solutions:

- Electrostatic longitudinal solution (i.e. $\tilde{E} // \vec{k}$). Longitudinal waves oscillate at $\omega = \omega_{pe}$ and are called plasma waves

- Electromagnetic transverse solutions: (i.e. $\vec{k} \cdot \vec{E} = 0$). In this case, the dispersion relation in the plasma reads:

$$|\vec{k}|^2 = \frac{\omega^2 - \omega_{pe}^2}{c^2} \quad (\text{III.11})$$

If $\omega > \omega_{pe}$, the electromagnetic wave propagates in the plasma with a wave vector

$$|k| = \sqrt{\frac{\omega^2 - \omega_{pe}^2}{c^2}}$$

On the other hand, if $\omega < \omega_{pe}$, the electromagnetic wave is evanescent within a scale length of the plasma skin depth $l_s \simeq \frac{c}{\omega_{pe}}$.

We can define the critical plasma density n_c associated to an electromagnetic wave frequency ω such as $\omega_{pe}(n_c) = \omega$.

$$n_c(\omega) = \frac{m_e \varepsilon_0 \omega^2}{e^2} \quad (\text{III.12})$$

A plasma medium is said to be **overcritical** if $n_e > n_c$. Otherwise, if $n_e < n_c$ the plasma is considered **underdense**.

When an electromagnetic wave, say a laser field, impinges on a plasma medium at normal incidence, it reflects off the plasma if the plasma density exceeds the laser critical density $n_e > n_c$, creating an evanescent wave at the plasma vacuum interface. At oblique incidence however, the reflection may occur even when the plasma density is below the critical density. Indeed, the Snell-Descartes law at the plasma vacuum interface reads:

$$\sin(\theta_0) = N(\omega) \sin(\theta_1) \quad (\text{III.13})$$

where θ_0 and θ_1 are the propagation angles of the electromagnetic field in the vacuum and the plasma respectively with respect to the plasma normal. Since $N(\omega) < 1$, the electromagnetic wave is deviated at the plasma surface such that $\theta_1 > \theta_0$. The maximum incidence angle θ_0 for which light still propagates inside the plasma is given by Snell-Descartes laws:

$$\begin{aligned} \sin(\theta_0) &= N(\omega) \\ &= \sqrt{1 - \frac{\omega_{pe}^2}{\omega^2}} \\ &= \sqrt{1 - \frac{n_e}{n_c}} \\ &\Leftrightarrow n_e = n_c \cos^2(\theta_0) \end{aligned} \quad (\text{III.14})$$

Therefore, the reflective density at oblique incidence reads:

$$n_{\text{ref}} = n_c \cos^2(\theta_0)$$

III.2 Gaussian laser formalism

Modelling intense laser pulses with plane waves is not satisfactory for HHG on plasma mirrors. Indeed, laser pulses are usually extremely focused on the PM surface and have an ultra-short temporal duration. For this matter, Gaussian beam formalism [27] constitutes a simplified analytical description of beams propagation, which is useful to develop models of laser-plasma interactions.

An electromagnetic laser field can be completely characterized by its transverse electric component $E(x, y, z, t)$ (assuming a transverse electromagnetic mode). This field verifies the scalar propagation equation:

$$\Delta E - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} E = 0 \quad (\text{III.15})$$

We assume that the laser pulse propagates along the z direction with a frequency ω_0 and we look for solutions of the form:

$$\begin{aligned} E(x, y, z, t) &= E_0(x, y, z, t) \exp\left(ik_0 z - i\omega_0 t\right) \\ k_0 &= \frac{\omega_0}{c} \end{aligned} \quad (\text{III.16})$$

Expanding the propagation equation, we find that E_0 verifies the following PDE:

$$\Delta E_0 - k_0^2 E_0 + 2ik_0 \frac{\partial E_0}{\partial z} - \frac{1}{c^2} \left[\frac{\partial^2 E_0}{\partial t^2} - \omega_0^2 E_0 - 2i\omega_0 \frac{\partial E_0}{\partial t} \right] = 0 \quad (\text{III.17})$$

Under the slowly varying envelope approximation, it is assumed that the electric field envelope amplitude E_0 slowly varies with respect to z and t variables. Formally, this approximation reads:

$$\begin{aligned} \left| \frac{\partial^2 E_0}{\partial z^2} \right| &\ll \left| k_0 \frac{\partial E_0}{\partial z} \right| \\ \left| \frac{\partial^2 E_0}{\partial t^2} \right| &\ll \left| \omega_0 \frac{\partial E_0}{\partial t} \right| \end{aligned} \quad (\text{III.18})$$

Therefore, the expanded propagation [III.17](#) equation could be simplified to obtain the parabolic wave equation:

$$k_0 \left(\frac{\partial E_0}{\partial z} + \frac{\partial E_0}{c \partial t} \right) - \frac{i}{2} \Delta_{\perp} E_0 = 0 \quad (\text{III.19})$$

with $\Delta_{\perp} = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ the transverse Laplace operator. To resolve equation [III.19](#), we look for solutions that can be expressed in the form:

$$E_0(x, y, z, t) = u(x, y, z) \times v(z, t)$$

$$\begin{aligned} k_0 \left(u'v + uv' + \frac{1}{c} u\dot{v} \right) - v \frac{i}{2} \Delta_{\perp} u &= 0 \\ \Leftrightarrow v \left(k_0 u' - \frac{i}{2} \Delta_{\perp} u \right) + k_0 u \left(v' + \frac{1}{c} \dot{v} \right) &= 0 \end{aligned} \quad (\text{III.20})$$

where the coma and upper dot operators denote the differentiation over the z and t variables respectively. $u(x, y, z)$ models the transverse spatial envelope of the electric field, while $v(z, t)$ models its temporal envelope. An admissible solution of equation III.20 is such that:

$$k_0 u' - \frac{i}{2} \Delta_{\perp} u = 0 \quad (\text{III.21})$$

and

$$v' + \frac{1}{c} \dot{v} = 0 \quad (\text{III.22})$$

The first equation III.21 represents the Helmholtz equation under the paraxial approximation. Admissible solutions of this equation include Gaussian beam solutions [27] (cf figure III.1).

$$u(x, y, z) = \frac{w_0}{w(z)} \exp \left[-\frac{x^2 + y^2}{w^2(z)} \right] \exp \left[-ik_0 \frac{x^2 + y^2}{2R(z)} \right] \quad (\text{III.23})$$

We define the terms involved in this equation as following:

- w_0 is called the laser waist. Under the slowly varying envelope approximation, the waist needs to be larger than the laser wavelength $w_0 > \lambda_0 = \frac{2\pi}{k_0}$.
- $w(z) = w_0 \sqrt{1 + (z/Z_r)^2}$ is the laser radius at which the laser intensity falls to $1/e^2$.
- $R(z) = z[1 + (Z_r/z)^2]$ the laser radius of curvature.
- with $Z_r = \frac{\pi w_0^2}{\lambda_0}$ the laser Rayleigh length. At $z = Z_r$, the laser radius grows by a factor of $\sqrt{2}$ and its radius of curvature is minimal.

We can also define the laser beam divergence:

$$\begin{aligned} \tan\theta &= \lim_{z \rightarrow +\infty} \frac{w(z)}{z} \\ \theta &\simeq \frac{\lambda_0}{\pi w_0} \quad (\text{under the paraxial approximation}) \end{aligned} \quad (\text{III.24})$$

The divergence of a Gaussian beam represents the angle covered by the laser cone of light far from the laser focus.

Equation III.22 is simply a homogeneous transport equation whose solution can be obtained by scattering the initial condition $v(\cdot, t = 0)$ over the characteristic lines ($z - ct = \text{constant}$). Thus, the solution of III.22 reads:

$$v(z, t) = v(z - ct, t = 0) \quad (\text{III.25})$$

The slowly varying envelope approximation requires the envelope width to be larger than the laser variation characteristic length (i.e. the laser wavelength in the space). A Gaussian temporal envelope is often employed to model v (but other representations are possible):

$$v(z, t) = E^{max} \exp \left(-\left[\frac{t - z/c}{\tau} \right]^2 \right) \quad (\text{III.26})$$

where τ is the temporal laser duration, with ($\tau \gg \frac{2\pi}{\omega_0} = T_0$) to respect the slowly varying envelope approximation. We can also define the laser intensity full-width at half maximum:

$$\text{FWHM} = \sqrt{2\ln(2)}\tau \simeq 1.177\tau$$

This quantity is more often used to characterize the temporal durations of ultra-short lasers.

The laser maximum amplitude E^{max} is an important parameter in laser-plasma interactions. We can define the laser normalized amplitude:

$$a_0 = \frac{eE^{max}}{m_e\omega_0c} = \lambda_0[\mu m]\sqrt{I_0[W.cm^{-2}]/(1.37e^{18})} \quad (\text{III.27})$$

The critical value $a_0 \simeq 1$ is generally considered as the threshold around which a laser pulse can drive relativistic electronic dynamics. In the context of laser-plasma interaction processes, the normalized laser amplitude is a key parameter that determines the relativistic or classical nature of the interaction.

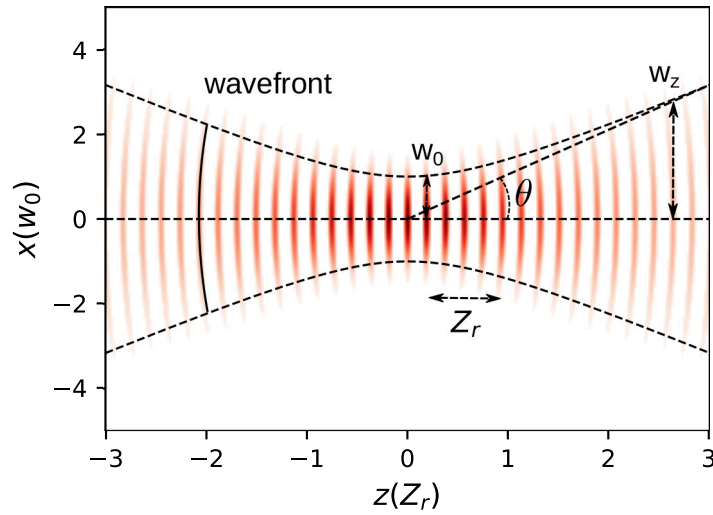


Figure III.1: Spatial profile of a Gaussian laser pulse (transverse electric field).

III.3 High harmonic generation on plasma mirrors

III.3.1 Plasma mirror Formation

A plasma mirror consists in an overdense plasma slab, created by irradiating an initially solid target with an intense laser pulse ($I > 10^{12}Wcm^{-2}$). In these conditions, atoms are partially or totally ionized, giving rise to an overdense warm plasma $n_e \gg n_c$. After the ionization process took place, and assuming that the plasma is isothermal, it undertakes thermal expansion in vacuum [28, 29] where the electrons thermal energy is transferred to ions. This expansion leads to the formation of an exponential plasma density gradient at the plasma vacuum interface as shown by figure III.2:

$$n_i(x, t) = n_0 \exp\left(\frac{x}{c_s t}\right) \quad (\text{III.28})$$

where c_s is the ion sound velocity given by:

$$c_s = \sqrt{\frac{Zk_B T_e}{m_i}} \quad (\text{III.29})$$

where k_B is the Boltzmann constant and T_e the electrons temperature.

The plasma density gradient characteristic length $L = c_s t$ is commonly called the plasma density scale length. For a 800 nm wavelength and a laser intensity of $I = 10^{16} \text{W.cm}^{-2}$, the plasma expands at a velocity of ~ 50 nanometers per picosecond. To guarantee a good plasma mirror optical quality, the PM density scale length should be much smaller than the laser wavelength $L \ll \lambda_0$ so that the laser interacts with a near-flat mirror before being reflected. This constraint shows that creating plasma mirrors requires lasers with a very high contrast. Indeed, typical UHI laser pulses are usually accompanied by a powerful laser pedestal ($I > 10^{12} \text{W.cm}^{-2}$) that impinges on the target tens of picoseconds before the main laser peak. This energy flux triggers a plasma vacuum expansion much more than a micrometer scale length before the main laser pulse reaches the target. To address this issue, a contrast improvement device, called DPM (for double-plasma mirror), is used to separate the main laser pulse from its pedestal, which results in improving the temporal contrast more than ten orders of magnitude but at the cost of an important loss of energy in the main pulse.

In practical HHG experiments, a plasma mirror is created by picking off a small fraction of the laser main pulse and focusing it before the laser main pulse. This way, the laser main pulse impinges on an already formed plasma mirror and the plasma density scale length can be controlled experimentally by adjusting the time delay between the main laser pulse and the prepulse [30, 31, 32, 33]. The plasma density scale length parameter has a strong impact over the physical mechanisms underlying the harmonic emission [34, 35, 36] and significantly affects the high order harmonic generation efficiency.

In the context of this thesis, we are mainly interested in the interaction between an ultra-intense, ultra-short laser pulse (i.e. the main laser pulse) with intermediate plasma expansion scale lengths (around $L \sim \frac{\lambda_0}{10}$). This regime of interaction maximizes the harmonic generation efficiency in the ultra-relativistic regime ($a_0 \gg 1$) [37].




Figure III.2: Schematic representation of plasma expansion into vacuum. Panel (a): initial solid target density profile. Panel(b): The formed plasma expands into vacuum to form an exponential plasma density gradient (the plasma expansion is exaggerated).

III.3.2 Harmonic generation mechanisms

The plasma mirror response to an ultra-intense laser illumination is governed by highly non-linear processes. For illustration, figure III.3 sketches the transverse electric field reflected by a plasma mirror irradiated by a p-polarized laser field for $a_0 = 10, \lambda_0 = 800 \text{ nm}, \theta = 55^\circ$ and $L = \frac{\lambda_0}{8}$ obtained from a 1D PIC simulation. The reflected field is periodically distorted with the same period as the incident laser field T_0 (panel (a)). The Fourier spectrum of the reflected field (panel (b)) exhibits a sequence of the spectral components $\omega_n = n\omega_0$ with $n \in \mathbb{N}^*$. These components are called high order harmonics. One can note that the high order harmonics span over tens of harmonic orders, reaching the X-UV emission regime ($3 \text{ nm} < \lambda < 100 \text{ nm}$). In panel (c), we filtered the harmonic signal between the 30th and 50th order. In the time domain, they form a train of extremely short light pulses, with a sub-femtosecond characteristic temporal duration each. This structure is called the *train of attosecond light pulses*.

f

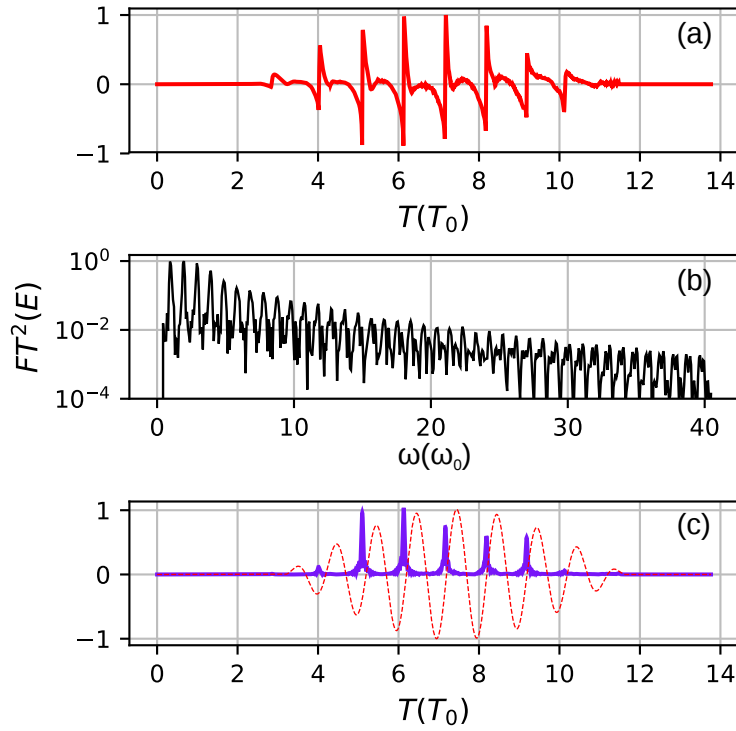


Figure III.3: Spectral and temporal response of a PM to an ultra-intense laser pulse obtained from a 1D PIC simulation ($a_0 = 10, L = \frac{\lambda_0}{8}, \theta_0 = 55^\circ$). Panel (a): Temporal evolution of the reflected field (Transverse magnetic field component). Panel (b): harmonic spectrum of the reflected field. Panel (c): filtered and normalized reflected fields associated to the laser frequency (red dashed curve) and harmonics between orders 30 and 50.

III.3.3 Attosecond pulses emission

This paragraph shows how a periodic spectrum in the Fourier space maps into a periodic train of extremely short light pulses in the temporal domain [38]. Let us consider a spectral profile $\tilde{S}(\omega)$ given by:

$$\tilde{S}(\omega) = \tilde{A}(\omega) \times [\tilde{H}(\omega) \otimes \tilde{E}(\omega)] \quad (\text{III.30})$$

where $\tilde{A}(\omega)$ is the spectrum envelope, $\tilde{H}(\omega)$ is a periodic Dirac comb with a period of ω_0 and $\tilde{E}(\omega)$ is the individual harmonic spectral profile. \times is the product operator \otimes is the convolution operator. The inverse Fourier transform of $\tilde{S}(\omega)$ reads:

$$S(t) = A(t) \otimes [H(t) \times E(t)] \quad (\text{III.31})$$

The temporal and spectral profiles of S , H and E are represented in figure III.4.

$A(t) = FT^{-1}(\tilde{A}(\omega))$ is an individual ultra-short (say attosecond) pulse with a duration of $\Delta T \propto \frac{1}{\Delta\omega}$ and where $\Delta\omega$ is the spectral bandwidth of \tilde{A} . The inverse Fourier transform of the Dirac comb $\tilde{H}(\omega)$ is a periodic Dirac comb with a period of $T_0 = \frac{2\pi}{\omega_0}$. The resulting temporal signal $S(t)$ is a periodic train (with a period of T_0) of attosecond light pulses, having a temporal envelope of $E(t) = FT^{-1}(\tilde{E}(\omega))$. Therefore, for a flat spectral phase, the total duration of the attosecond pulses train is associated to the individual harmonic profiles, while the temporal duration of an individual attosecond pulse is pertaining to the spectral width of the whole harmonic spectrum.

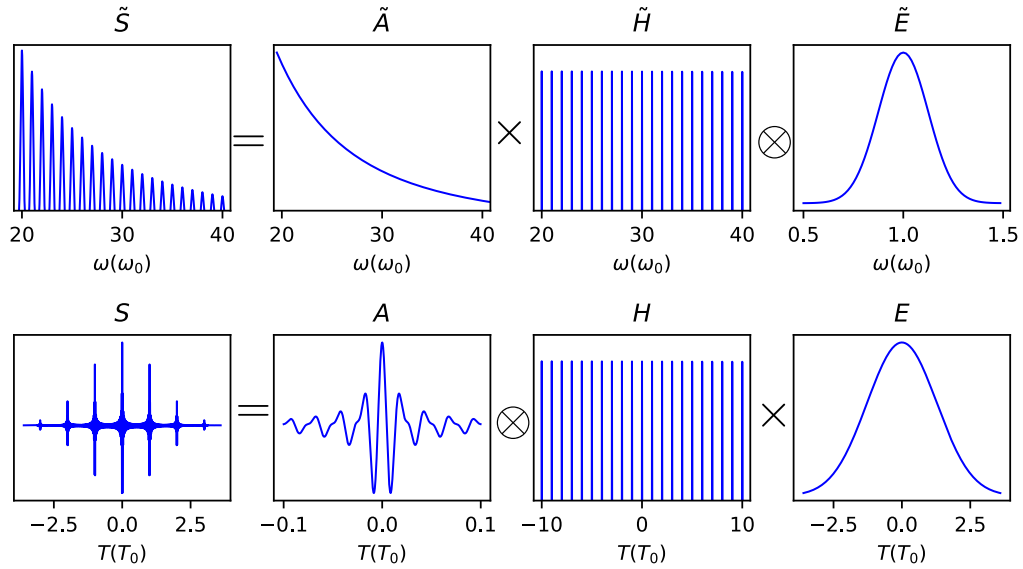


Figure III.4: Harmonic spectrum pertaining to an attosecond pulses train. Upper and lower panels are related by Fourier transforms.

It is important to mention that deducing the temporal durations of the attosecond pulses train or of the individual pulses from the harmonic spectrum requires a known spectral phase $\phi(\omega) = \arg(\tilde{S}(\omega))$. Obtaining a perfectly periodic attosecond pulses train with the shortest attosecond pulses durations (i.e. Fourier limited durations) is achieved with constant phase across all harmonic orders.

Femtosecond chirp

An example of the potential impact of a non-constant spectral phase is discussed in this paragraph. Consider a harmonic signal $S(t)$ with a Gaussian temporal envelope $E(t) = \exp\left(\frac{-t^2}{\tau^2}\right)$. In this case,

$$S(t) = E(t) \sum_n A_n \exp[in\omega_0 t]$$

We now introduce a temporal phase $\phi(t) = n\alpha t^2$ to each harmonic order:

$$\begin{aligned} S(t) &= E(t) \sum_n A_n \exp[int\omega_0 + i\phi_n(t)] \\ S(t) &= E(t) \sum_n A_n \exp[int(\omega_0 + \alpha t)] \end{aligned} \quad (\text{III.32})$$

In the presence of non-constant harmonic phase, the train periodicity is distorted (the temporal delay between successive attosecond pulses is no longer equal to T_0) as shown by figure III.5.

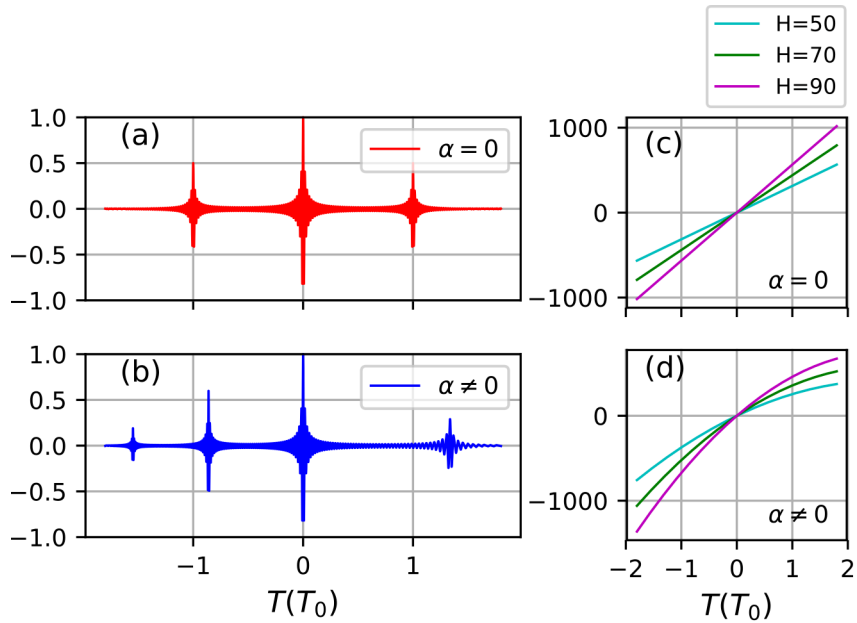


Figure III.5: Attosecond pulses train with (panel (b)) and without (panel (a)) femtosecond chirp. Panels (c) and (d) sketch the evolution of the temporal phases of three harmonic orders (50,70 and 90) for both cases.

III.4 Models for HHG on plasma mirrors

III.4.1 Bourdier boosted frame

Analytically modelling HHG on plasma mirrors can be extremely complicated, especially when taking into account the full dimensionality of the problem. Although, if the incident laser pulse is approximated by a plane-wave, Bourdier [39] showed that it is possible to reduce this problem dimensionality to 1D via Lorentz frame transformation. This operation is illustrated in figure III.6.

We consider a plasma mirror occupying the half space $x < 0$ in the laboratory frame \mathcal{L} . A laser beam impinges the plasma with an angle of incidence θ in the (x, z) plane. The laser wave vector

reads:

$$\vec{k}_0^{\mathcal{L}} = \sin(\theta)\vec{e}_z - \cos(\theta)\vec{e}_x$$

The Lorentz Bourdier boosted frame \mathcal{L}' moves with a velocity $v = -c\sin(\theta)\vec{e}_z$ with respect to the laboratory frame \mathcal{L} . The relativistic factors of this Lorentz transformation are $\beta_0 = \sin(\theta)$ and $\gamma_0 = \frac{1}{\sqrt{1-\beta_0^2}} = \frac{1}{\cos(\theta)}$. The laser wave vector is:

$$\vec{k}_0^{\mathcal{L}'} = -\frac{\|\vec{k}_0^{\mathcal{L}}\|}{\gamma_0}\vec{e}_x = -\|\vec{k}_0^{\mathcal{L}}\|\cos(\theta)\vec{e}_z$$

In the Bourdier frame the laser is normal to the target and the laser frequency is decreased by a factor γ_0 :

$$\omega_0^{L'} = \frac{\omega_0}{\gamma_0}$$

On the other hand, the plasma is no longer at rest. Instead, it drifts along the z-axis with a velocity of $\sin(\theta)c\vec{e}_z$. Due to the length contraction induced by the Lorentz transformation along the z-axis ($\delta z^{L'} = \frac{\delta z^L}{\gamma_0}$), the plasma density is increased by a factor of γ_0 :

$$n^{\mathcal{L}'} = \gamma_0 n^{\mathcal{L}}$$

The Lorentz transformation over the transverse electric and magnetic fields, assuming a p-polarized laser pulse (i.e. the magnetic field is orthogonal to the incidence plane) reads:

$$\begin{aligned} \vec{E}_t^{\mathcal{L}'} &= \frac{E_0}{\gamma_0}(\vec{e}_z) \\ \vec{B}_t^{\mathcal{L}'} &= \frac{E_0}{c\gamma_0}(\cos(\theta)\vec{e}_y) \end{aligned} \quad (\text{III.33})$$

The laser field amplitude is reduced by a factor of γ_0 due to the Lorentz transform. However, the laser normalized amplitude a_0 remains unchanged.

The Bourdier boosted frame is a powerful tool that enables both simpler analytical and numerical modellings of laser-plasma interaction at oblique incidence. In chapter VI, we also show how to employ Lorentz transforms as an efficient diagnostic tool of laser-plasma interaction simulations for arbitrary electromagnetic laser waves (and not just plane waves).




Figure III.6: Illustrative scheme for laser-plasma interaction in the laboratory frame (left) and the Bourdier boosted frame (right).

III.4.2 Wave equation with sources in Bourdier frame

In this subsection, we derive the field radiated by the plasma currents in the Bourdier frame. We get rid of the \mathcal{L}' specifying the Bourdier frame. We describe the electromagnetic fields using scalar and vector potentials (V, \vec{A}) under the Coulomb gauge:

$$\begin{aligned}\vec{E} &= -\vec{\nabla}V - \frac{\partial \vec{A}}{\partial t} \\ \vec{B} &= \vec{\nabla} \wedge \vec{A} \\ \vec{\nabla} \cdot \vec{A} &= 0\end{aligned}\tag{III.34}$$

The radiated fields are determined by the vector potential \vec{A} satisfying [40]:

$$\Delta \vec{A} - \frac{\partial^2 \vec{A}}{\partial t^2} = \mu_0 \vec{J}_\perp\tag{III.35}$$

where $\vec{J}_\perp = (0, J_y, J_z)$ is the transverse current. This system is invariant along both transverse directions y and z, so the Laplacian operator can be simplified to a second order derivative along x-axis $\Delta = \frac{\partial^2}{\partial x^2}$. To integrate equation III.35 under these symmetry conditions, we employ the Green function $G(x, t)$ defined by:

$$G(x, t) = \Theta\left(t - \frac{|x|}{c}\right)\tag{III.36}$$

solution of:

$$\left(\frac{\partial^2}{\partial z^2} - \frac{1}{c^2} \frac{\partial^2}{\partial t^2}\right) G(x, x', t, t') = -\mu_0 \delta_0(x - x') \delta_0(t - t')\tag{III.37}$$

where Θ and δ_0 respectively denote the Heaviside function and the Dirac pulse.

The radiated field by the plasma is:

$$\begin{aligned}\vec{A}^r(x_0, t_0) &= -c\mu_0 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G(x_0 - x, t_0 - t) \vec{J}_\perp(x, t) dt dx \\ \vec{A}^r(x_0, t_0) &= -c\mu_0 \int_{-\infty}^{\infty} \int_{-\infty}^{t_0 - |x_0 - x|/c} \vec{J}_\perp(x, t) dt dx\end{aligned}\tag{III.38}$$

Note that the total electromagnetic field (radiated + incident) is:

$$\vec{A}(x_0, t_0) = \vec{A}^r(x_0, t_0) + \vec{A}^0(x_0, t_0)\tag{III.39}$$

where $\vec{A}^0(x_0, t_0)$ is the incident laser field which verifies equation III.35 without the right hand side. Therefore, the transverse electric field radiated by the plasma is determined by:

$$\vec{E}_\perp^r(x_0, t_0) = -\frac{\partial \vec{A}^r}{\partial t} = -c\mu_0 \int_{-\infty}^{\infty} \vec{J}_\perp(x, t_0 - |x_0 - x|/c) dx\tag{III.40}$$

Equation III.38 shows that the radiated electromagnetic field can be rigorously determined provided the transverse currents in the plasma.

III.4.3 Transverse currents expression

We now derive the transverse currents expression in the Bourdier frame. Considering the plasma medium as a fluid of electrons and ions, the transverse current is:

$$\vec{J}_\perp = -en_e\vec{v}_{e,\perp} + Zen_i\vec{v}_{i,\perp} \quad (\text{III.41})$$

where $\vec{v}_{e,\perp}$ and $\vec{v}_{i,\perp}$ are the electronic and ionic transverse velocities. Assuming immobile ions on the laser duration time scale (due to their much higher mass), their velocity in the Bourdier frame remains equal to the plasma drift velocity $\vec{v}_{i,\perp} = c\sin(\theta)\vec{e}_z$.

On the other hand, electron contribution to the current is obtained by using the conservation of the transverse canonical momentum, which results from the translational invariance of the system along transverse direction y and z.

$$\begin{aligned} \frac{d(\vec{p}_\perp - e\vec{A})}{dt} &= 0 \\ \Leftrightarrow \vec{p}_\perp &= \vec{p}_\perp(t=0) + e\vec{A} \end{aligned} \quad (\text{III.42})$$

where

$$\vec{p}_\perp(t=0) = -m_e c \tan(\theta)\vec{e}_z$$

is the electron momentum before the arrival of the laser.

Finally, the total transverse current in Bourdier frame reads:

$$\begin{aligned} \vec{J}_\perp(x,t) &= -\frac{e^2 n_e(x,t)\vec{A}(x,t)}{m_e \gamma(x,t)} - ec \tan(\theta) \left(Zn_i(x,t)\cos(\theta) - \frac{n_e(x,t)}{\gamma(x,t)} \right) \vec{e}_z \\ \gamma(x,t) &= \sqrt{\frac{1 + \|e\vec{A}/(cm_e) - \tan(\theta)\vec{e}_z\|^2}{1 - (v_x^e/c)^2}} \end{aligned} \quad (\text{III.43})$$

The first equation of [III.43](#) shows that the transverse electric current in the plasma is constituted of two terms. The first term, which is $\propto \frac{n_e}{\gamma}\vec{A}$, corresponds to the current driven by the radiation field \vec{A} . The second term only appears whenever the electronic and ionic drift currents do not compensate one another and vanishes at normal incidence. This can happen if the electronic and ionic densities ($n_e \neq n_i$) are no longer equal, or if the electrons get accelerated by the laser field ($\gamma \neq \frac{1}{\cos(\theta)}$). This current is always oriented along z-axis, and is independent from the laser polarization. This shows that a p-polarized harmonic signal can be emitted by the PM even if the laser is s-polarized (i.e, if the laser electric field is normal to the plane of incidence).

According to equation [III.43](#), additional frequencies in the transverse electric current could emerge from electronic density distribution or Lorentz factor modulations. Density fluctuations play a crucial role in HHG through the coherent wake emission process (CWE) [[16](#), [41](#), [42](#)].

This analysis is however not to determine the spectrum of the radiated light, which is also sensitive to the spatial variation of the transverse current \vec{J}_\perp through the integral equation [III.40](#). In particular,

in the relativistic regime, the plasma mirror surface longitudinal oscillation can result in a harmonic radiation through Doppler effect [12, 43].

III.4.4 Coherent wake emission regime

For laser intensities not exceeding the relativistic threshold $a_0 \leq 1$ (see equation III.27), and very steep plasma density scale lengths $L \leq \frac{\lambda_0}{20}$, harmonic generation on plasma mirrors is dominated by the coherent wake emission mechanism [16].

This mechanism is illustrated in figure III.7. It results from the excitation of plasma waves in the plasma density gradient by Brunel electrons [44], sent back in the overdense part of the plasma gradient under the combined effects of the charge separation forces and the laser electric field. When traveling inside the plasma gradient, these electrons trajectories cross each other (because the kinetic energy differs from one electron to another), forming a density electrons peak that excites plasma waves at the local plasma frequencies. Hence, all the plasma frequencies between ω_0 and ω_{pe} are excited, resulting in a coherent emission (because the plasma density scale length is very short compared to λ_0) of harmonic radiations along the specular direction through linear mode conversion [45]. This mechanism takes place at each laser optical cycle, resulting in a train of attosecond light pulses associated to a periodic harmonic spectrum.

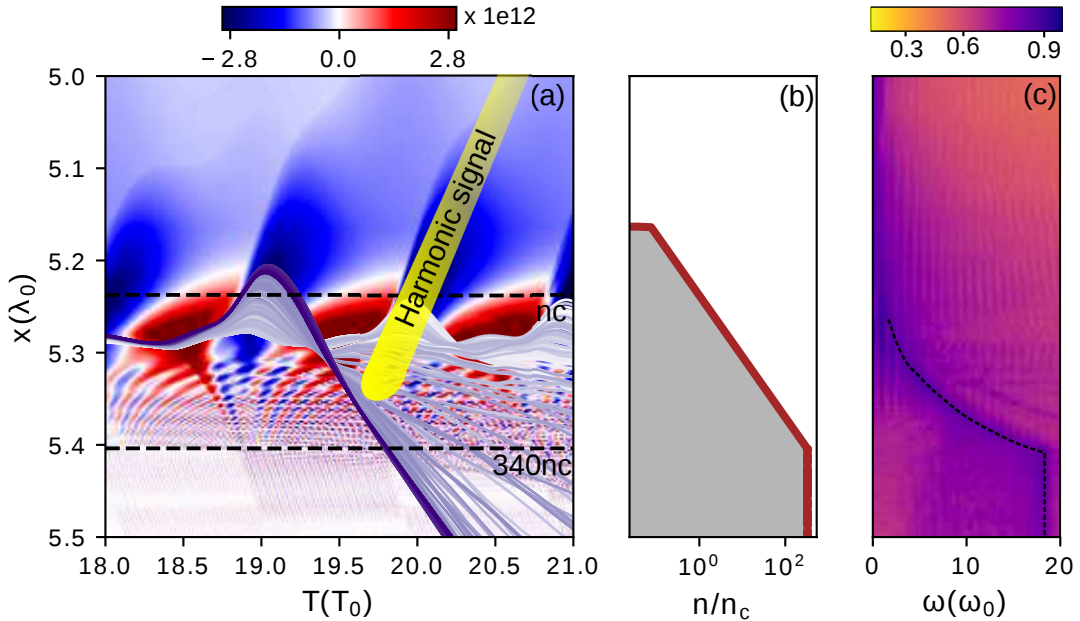


Figure III.7: Illustration of the CWE mechanism. The blue-white-red colormap in panel (a) sketches the longitudinal plasma waves (E_x field) at the surface of the plasma mirror obtained from a 1D CWE PIC simulation (with $a_0 = 0.6$, $L = \frac{\lambda_0}{20}$, $\theta_0 = 55^\circ$, $n_{max} = 340n_c$). The purple curves represent the Brunel electrons trajectories, intersecting inside the plasma. Panel (b) illustrates the initial plasma density in units of n_c . Panel (c) is the Fourier transform over time of the plasma waves excited inside the plasma. Note that the frequencies at each depth match the plasma frequencies associated to the local plasma densities.

Note that the highest excited plasma frequency is the frequency of the plasma bulk ω_{pe}^{max} . Conse-

quently, the CWE mechanism is characterized by a harmonic cutoff given by:

$$\omega_{cwe}^{\max} = \omega_{pe}^{\max} = \sqrt{\frac{n_{\max}}{n_c}} \omega_0$$

The laser envelope spatial or temporal variations critically affects the electronic peak velocity inside the plasma gradient. Indeed, for a stronger laser electric field, the Brunel electrons gain more kinetic energy from the laser before traveling inside the plasma. Therefore, the trajectories intersections, as well as the harmonic emission, occur earlier at each laser optical cycle. Thus, CWE attosecond pulses train is not perfectly periodic (i.e. femtosecond chirp) because of the laser envelope temporal variation. In the same way, CWE harmonics wavefronts are concave because of the laser intensity variation across the focal spot [42].

III.4.5 Harmonic generation in the relativistic regime

For ultra-high intensity lasers ($I \geq 10^{19} \text{Wcm}^{-2}$), corresponding to a normalized vector potential amplitude $a_0 \geq 1$, the PM surface acts as a relativistic oscillating mirror that simultaneously reflects and compresses the incoming laser field, giving rise to additional high order harmonics. This effect has been studied, both theoretically and experimentally, for decades [12, 43, 46, 15]. It constitutes a promising path towards achieving extremely short (attosecond), intense XUV sources. In this section, we first review the principle of the Doppler effect induced by a mirror at a constant speed and present the Lichters analytical approach to explain the Doppler harmonic generation process.

III.4.6 Doppler effect

Before describing the Doppler effect induced by an oscillating mirror, we first recall the frequency shift induced by a perfectly reflective mirror following a uniform motion with a velocity $v_0 \vec{e}_x$.

Consider a monochromatic wave with a frequency of ω that is specularly reflected by the moving mirror (cf figure III.8). A receiver, at rest in the laboratory frame observes a reflected wave with a frequency [47]:

$$\omega_r = (1 + \beta)^2 \gamma^2 \omega$$

where $\beta = v_0/c$ and $\gamma = \frac{1}{\sqrt{1-\beta^2}}$. The frequency of the reflected field increases when the reflective mirror moves towards the observer ($v_0 > 0$). This frequency shift is called the relativistic Doppler effect. This shift increases with the mirror velocity. When the mirror motion is ultra-relativistic ($\gamma \gg 1$), the Doppler shift reaches $\simeq 4\gamma^2$.




Figure III.8: Doppler shift induced by a reflective mirror moving at constant speed in the specular direction.

III.4.7 The Relativistic Oscillating Mirror (ROM) model

In 1996, R. Lichters *et al* [12] presented a theoretical model (the ROM model) to account for the high order harmonic generation occurring in the relativistic regime.

The radiated electric field is related to the transverse electric current by equation III.40. Lichters *et al* assume that the radiative electrons are located within the skin depth at the plasma mirror surface. Moreover, they assume that these electrons undergo a fast longitudinal oscillations under the effect of the laser electric field component normal to the plasma surface. As the plasma skin depth is very small compared to the laser wavelength ($l_s = \frac{c}{\omega_{pe}} \ll \lambda_0$), it is reasonable to model the radiative electrons as a punctual source at a position $Z(t)$ that varies as:

$$Z(t) = \frac{c}{\omega_0} \frac{2A_0(x, t)\sin(\theta)}{1 + 2A_0(x, t)\sin(\theta)} \cos(\omega_0 t) \quad (\text{III.44})$$

where A_0 is the temporal envelope of the laser vector potential \vec{A}_i . The electric current is determined by equation III.43. It is supposed to vanish over a scale length of l_s beyond $Z(t)$. Consequently, equation III.40 simplifies to:

$$E_r(z, t) \simeq l_s c \mu_0 J_{\perp}(Z(t_{ret}), t_{ret}) \quad (\text{III.45})$$

where t_{ret} is the retarded time at which the signal observed by an observer located at position z and time t has been emitted by the PM. This means that:

$$t_{ret} = t - \frac{(Z(t_{ret}) - z)}{c} \quad (\text{III.46})$$

In case of a p-polarized laser field, the radiated electric field reads:

$$\begin{aligned} \vec{E}_r(z, t) = \vec{E}_r(z, t) = \frac{\omega_p}{2\omega_0} \left[\frac{\sqrt{1 - (Z(t_{ret})/c)^2} (A_0(Z(t_{ret}), t_{ret}) - \tan(\theta))}{\sqrt{1 + A_0^2(Z(t_{ret}), t_{ret}) \cos^2(\theta) - \sin(2\theta) A_0(Z(t_{ret}), t_{ret})}} \right. \\ \left. + \tan(\theta) \left(1 + Z(t_{ret})/l_s \right) \right] \vec{e}_x \end{aligned} \quad (\text{III.47})$$

To determine t_{ret} and therefore $Z(t_{ret})$, Lichters *et al* propose to use a fixed point iterative method:

$$\begin{aligned} Z_0 &= Z(t) \\ Z_{n+1} &= Z\left(t - \frac{Z_n - z}{c}\right) \end{aligned} \quad (\text{III.48})$$

The Z_n sequence numerically converges to $Z(t_{ret})$ and $t_{ret} = t - \frac{Z(t_{ret}) - z}{c}$.

Figure III.9 illustrates numerical results obtained from the ROM model in the case of an obliquely incident laser pulse at 55 degrees and $a_0 = 6$. The laser temporal envelope is Gaussian with a duration of $\tau = 2T_0$. Panel (a) sketches the laser electric field impinging on the PM. Panel (b) illustrates the plasma mirror surface dynamics $Z(t)$ (blue curve) and the retarded dynamic $Z(t_{ret})$ (red curve). One can note that $Z(t_{ret})$ clearly differs from a pure trigonometric function. This gives rise to a strongly distorted reflected field (panel(c)). This temporal distortion is associated with the presence of high order laser harmonics in the spectral domain and attosecond light pulses in the temporal domain.

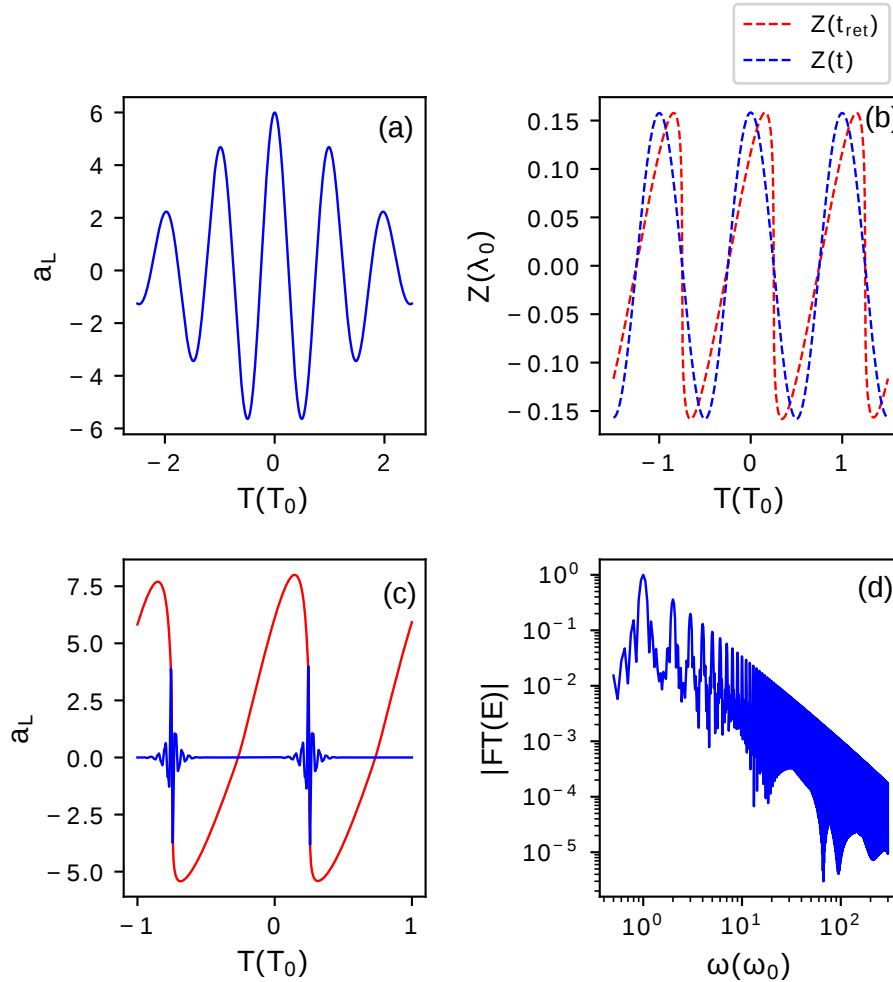


Figure III.9: High order harmonic generation with the ROM model. Panel(a): Incident laser profile. Panel(b): plasma surface dynamic (retarded: red, unretarded: blue). Panel (c): Reflected field (red) and attosecond pulses associated to harmonic orders between 20 and 60 (blue). Panel (d): Spectral profile of the reflected field in log scale.

The ROM model is a simple and intuitive approach to qualitatively understand the high order harmonic generation process on overdense plasmas in the relativistic regime. However, it largely omits the physics at play at the PM surface and therefore lacks of predictive power regarding harmonic generation scaling laws and generation efficiency. Moreover, imposing a predefined dynamics to the PM surface is not satisfactory. Figure III.10 sketches the PM dynamics and the HHG emission obtained from 1D PIC simulation (in the Bourdier frame) in the relativistic regime ($a_0 = 10$, $L = \frac{\lambda_0}{8}$, $\theta = 55^\circ$). Even though a single high order harmonic bunch (in the form of an attosecond light pulse) is emitted at each laser optical cycle, the PM dynamics is more unpredictable and does not follow a single order harmonic motion.




Figure III.10: Plasma mirror dynamics and HHG for $a_0 = 10$ at 55° incidence angle from 1D PIC simulation (in the Bourdier frame). In gray, the electronic density at the plasma-vacuum edge. In red, high order harmonics in the range 20-30

Moreover, this modelling completely omitted multiple aspects that can only be apprehended in higher dimensions, such as the spatial properties of Doppler harmonics. In the next section, we present a comprehensive model developed in H.Vincenti's thesis [48] [49] to account for the spatial properties of Doppler harmonics. This model investigates the plasma mirror surface denting in the presence of a plasma density ramp. Based on this dynamic, it is then possible to derive an accurate predictive model for Doppler harmonic spatial properties.

III.4.8 Plasma denting model and spatial properties of Doppler harmonics

In the relativistic regime, the mean positions of both ions and electrons from the surface are pushed inwards due to the strong laser radiation pressure, resulting in a plasma mirror surface denting. This process can be described with a three-steps quasi-static model. First, assuming immobile ions, the maximum electrons excursion of the PM surface is modelled by determining the balance between the

pushing force exerted by the incident and reflected laser field, and the restoring force due to the space-charge separation. Second, the ion motion induced by the laser radiation pressure is computed using the momentum conservation. Finally, this ion dynamics is taken into account to update electrons excursion. This model assumes that the electronic motion can be computed for a fixed ion background each laser optical cycle, which is valid as long as the ion motion can be neglected during a single optical cycle.

Electrons excursion

We suppose an immobile ions background within the time scale of one laser period. In Bourdier frame (the plasma initially occupies the domain $x > 0$ and drifts along the z -axis), the laser transverse magnetic field B^y gives rise to a longitudinal force exerted on the transversally drifting electrons. As a result, the surface electrons are piled up to form a dense electronic spike that undergoes an excursion inside the plasma. We denote by $x_e(t)$ the longitudinal position of the electronic spike. The space charge separation simultaneously gives rise to two additional forces exerted on the electrons spike.

First, the electrostatic space-charge separation forces associated to the longitudinal field that reads:

$$E_x(x_e(t)) = \int_{-\infty}^{x_e(t)} Z n_i(x) dx$$

On the other hand, the transversally drifting ions induce an uncompensated electric current $J_z(x) = -Z n_i(x) c \sin(\theta)$ that gives rise to a magnetic field $B_y^d = -\mu_0 \int_{-\infty}^{x_e(t)} J_z(x) dx$.

When the maximum electronic excursion x_m is reached at a time t_e , we assume that an equilibrium between the three forces exerted on the electronic spike is achieved. Therefore, we have:

$$v_x [B_y^{laser}(z_m) + B_y^d(z_m)] + E_z(z_m) = 0 \quad (\text{III.49})$$

In the relativistic regime, we can assume that $v_x \simeq c$ during the pushing phase. For an exponential plasma density ramp with a scale length L , we get a maximum plasma excursion.

$$x_m = L \ln \left[1 + \frac{2a_0(1 + \sin(\theta))}{2\pi L/\lambda_0} \frac{n_c}{n_e(x_i)} \right] \quad (\text{III.50})$$

The maximum electronic excursion increases with both the laser amplitude or the plasma density scale length: an increase in the laser intensity directly increases the laser force pushing the electronic spike. On the other hand, increasing the plasma density scale length reduces the space charge separation forces due to the decrease in the density in the plasma ramp.

Ions excursion

Under the effect of space-charge separation forces, ions start to slowly travel towards the electronic spike inside the plasma. Applying momentum/energy conservation, one can be shown that the ions excursion in the plasma can be expressed as:

$$x_i(t) = 2L \ln \left[1 + \frac{c\Pi_0}{2L\cos(\theta)} \int_0^t a_L(t_0) dt_0 \right] \quad (\text{III.51})$$

where $\Pi_0 = \sqrt{\frac{RZm_e \cos(\theta)}{2Am_p}}$ with R the laser reflexion coefficient, A the ions mass number, and $a_L(t)$ the laser vector potential normalized temporal envelope. Similarly to electrons, the ions excursion increases with the laser intensity and the plasma density scale length.

Plasma mirror excursion

The total plasma excursion is simply given by:

$$x_T(t) = x_i(t) + x_e(t) \quad (\text{III.52})$$

In the expression of $x_e(t)$ (equation III.50), the term $\frac{n_e}{n_e(x_i)}$ is equal to $\exp(-x_i/L)$:

$$x_e(t) = L \ln \left[1 + \frac{2a_L(t)(1 - \sin(\theta))}{2\pi L/\lambda_0} \exp[-x_i(t)/L] \right] \quad (\text{III.53})$$

Therefore, III.52 reads:

$$x_T(t) = x_i(t) + L \ln \left[1 + \frac{2a_L(t)(1 - \sin(\theta))}{2\pi L/\lambda_0} \exp[-x_i(t)/L] \right] \quad (\text{III.54})$$

In the relativistic regime, the laser radiation pressure tends to dent the PM surface. In return, the deformation of the plasma surface affects the spatio-temporal properties of the radiated harmonic beam. As the density spikes responsible for the HHG undergo a deeper inward excursion from one laser optical cycle to another, the spacing between consecutive attosecond light pulses is expected to increase over time, which results in a femtosecond chirp of the harmonic signal. The impact of the plasma mirror deformation on the spatial properties of the harmonic signal is detailed in the next paragraph.

III.4.9 Plasma mirror excursion in three dimensions:

In higher dimensions (2D or 3D), the laser intensity variation along its cross section induces an inhomogeneous denting over the PM surface. Note that the Lorentz transform from/to the Bourdier frame only modifies the scale lengths along the transformed axis. Thus, the plasma excursion in the Bourdier and the lab frames are equal. Therefore, at each position (z, y) of the plasma surface (considering that the z - y plane is transverse to the target), the total denting reads:

$$\begin{aligned} X_T(t, z, y) &= x_i(t, z, y) + L \ln \left[1 + \frac{2a_L(x, y, t)(1 + \sin(\theta))}{2\pi L/\lambda_0} \frac{n_c}{n_0 \exp[-x_i(t)/L]} \right] \\ x_i(t, z, y) &= 2L \ln \left[1 + \frac{c\Pi_0}{2L \cos(\theta)} \int_0^t a_L(z, y, t_0) dt_0 \right] \\ a_L(z, y, t) &= a_0 \exp \left[-\frac{t^2}{\tau^2} \right] \exp \left[-\frac{(z \cos(\theta))^2 + y^2}{w_0^2} \right] \end{aligned} \quad (\text{III.55})$$

For $x, y \ll w_0$, we can use Taylor expansion to simplify equation III.55:

$$X_T(t, z, y) = X_0 - \frac{z^2 \cos^2(\theta) + y^2}{4f_p(t)} + o(x^2) + o(y^2) \quad (\text{III.56})$$

were X_0 is a constant, and:

$$\begin{aligned} f_p(t) &= \frac{w_0^2}{4L} \frac{\epsilon(t) + (1 + \mu(t)\epsilon(t))^2}{\epsilon(t) + 2\mu(t)\epsilon(t)(1 + \mu(t)\epsilon(t))} \\ \epsilon(t) &= \frac{2\lambda_0 a(t)}{2\pi L(1 - \sin(\theta))} \\ \mu(t) &= \frac{\omega_0(1 - \sin(\theta))\Pi_0}{4} \int_0^t \frac{a(t_0)}{a_0} dt_0 \end{aligned} \quad (\text{III.57})$$

Equation III.56 describes a paraboloid (cf a parabola in 3D) with focal distances $\frac{f_p}{\cos^2(\theta)}$ and f_p in the (x-z) and (x-y) planes respectively.

$\epsilon(t)$ and $\mu(t)$ are associated with the electronic and ionic excursions respectively. f_p is called the PM focal distance. We will show later that this physical quantity is of fundamental importance to determine the spatial properties of high order harmonics.

At the beginning of the laser-plasma interaction (the first laser optical cycles), the ionic contribution to the total plasma denting can be neglected ($\mu(t) \ll 1$) compared to the electronic motion. For ultra-relativistic laser intensities ($a_0 \gg 1$), the total plasma mirror dynamic is dominated by the electronic motion ($\epsilon(t) \gg 1$). Therefore III.57 simplifies to:

$$f_p \simeq \frac{w_0^2}{4L} \quad (\text{III.58})$$

Note that for very long laser pulses ($\tau \gg T_0$), and $t \gg T_0$ (after many optical cycles), the plasma mirror dynamics becomes more and more dominated by the ionic excursion and we have:

$$f_p \simeq \frac{w_0^2}{8L} \quad (\text{III.59})$$

We can also define the plasma mirror denting parameter by the difference between the plasma denting at the center of the interaction region and at a distance of $\sqrt{2}w_0$ in the incidence plane.

$$\delta_p = \frac{2w_0^2 \cos^2(\theta)}{f_p} \quad (\text{III.60})$$

This quantity relates the PM maximum denting to its focal distance f_p . Due to the spatial variations of the laser amplitude at focus, equation III.56 shows that the plasma target gets curved, therefore focusing the high order harmonics. In the following, we derive the spatial properties of these harmonics as a function of the PM curvature.

III.4.10 Spatial properties of high order harmonics

The plasma mirror denting takes the form of a parabolic curved surface, described by equation III.56. It can be shown that the focal distance of the PM slowly varies from one optical cycle to another [48]. During the first laser optical cycles, we assume that the focal distance of the PM is constant and is given by III.58. Therefore, the harmonic beam is emitted along the specular direction given by $-\vec{k}_r = k_0 \cos(\theta) \vec{e}_x + k_0 \sin(\theta) \vec{e}_z$ by a parabolic surface. We neglect the high order optical aberrations (e.g. coma aberration) induced by an off-axis reflection on a parabolic surface [50]. In the

reflection plane (see figure III.11), the plasma mirror parabolic curvature induces a quadratic spatial phase on the emitted harmonic beam. This quadratic phase originates from the cumulated optical path difference of different light rays emitted by the curved surface and it varies as:

$$\phi_p(x', y') = -2k_n \cos(\theta) Z(x', y') \quad (\text{III.61})$$

where k_n is the wave vector of the n^{th} harmonic order and (x', y') is the coordinate system of the emission plane:

$$\begin{aligned} x' &= x \cos(\theta) \\ y' &= y \end{aligned} \quad (\text{III.62})$$

Therefore, the spatial phase III.61:

$$\phi_p(x', y') = -k_n \frac{x'^2 + y'^2}{2f_p / \cos(\theta)} \quad (\text{III.63})$$

where k_n is the wave vector of the n^{th} harmonic order and (x', y') is the coordinate system of the emission plane. This equation describes a curved wavefront. The plasma mirror acts as a focusing optics on both the incident laser and the high order harmonics. Note that despite the oblique incidence in the (x, z) plane, the plasma mirror induces no astigmatism into the reflected field: the wavefronts are equivalently curved along the x' and y' directions. Therefore, the harmonic beam focuses at a distance $d = \frac{f_p}{\cos(\theta)}$ along the specular direction from the PM. This particular behavior suggests that the PM could be used as a focusing optics in order to reach extremely high electromagnetic intensities [37]. In this thesis, we will also show a configuration where the harmonic beam is astigmatic due to different wave-front curvatures along the two transverse directions.

In order to establish equation III.61, we assumed that the laser spatial phase is constant over the laser-plasma interaction region. This assumption holds true if the laser focus coincides with the PM surface. However, it might be very beneficial to use an incident laser out of focus (with curved wavefronts on PM) in various application experiments employing Doppler HHG. The general harmonic phase formula for an arbitrary laser wavefront curvature writes:

$$\begin{aligned} \phi_n(x', y') &= \phi_p + \phi_L \\ \phi_n(x', y') &= -k_n \left(\frac{x'^2 + y'^2}{2f_p / \cos(\theta)} - \frac{x^2}{2R_z} \right) \end{aligned} \quad (\text{III.64})$$

with R_z the laser radius of curvature at the PM emission plane.

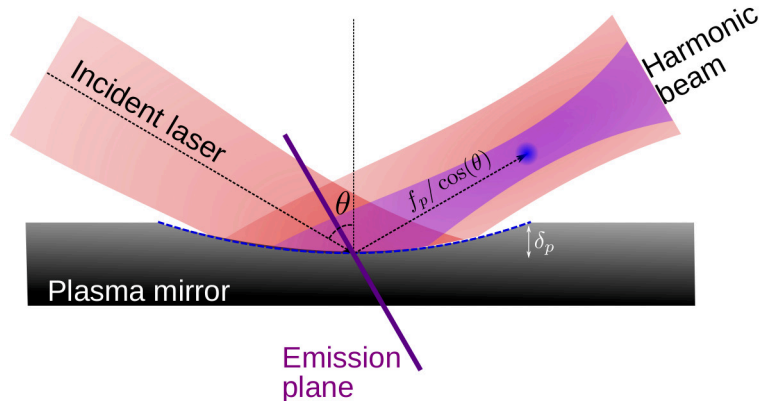


Figure III.11: Schematic representation of the plasma mirror denting inducing harmonic focusing.

Harmonic divergence

The divergence of the harmonic beam can be computed from the harmonic spatial phase given by equation III.63. It reads:

$$\theta_n = \lim_{z \rightarrow +\infty} \frac{w_n(z)}{z} \quad (\text{III.65})$$

where z represents the distance from the PM mirror along the specular reflection direction and $w_n(z)$ is the harmonic beam size at a distance z from the PM focus. Assuming that each individual harmonic can be modelled as a Gaussian pulse, we can use the Gaussian beam theory to determine $w_n(z)$. First, we can define the complex harmonic beam parameter q_n at the PM surface:

$$q_n(z - z_{rn}) = (z - z_{rn}) + iZ_r^n \quad (\text{III.66})$$

where z_{rn} is the distance between the PM emission plane and the focus of the n^{th} harmonic, and Z_r^n is the Rayleigh length of the same harmonic order, which writes :

$$Z_r^n = \frac{n\pi w_n^2}{\lambda_0} \quad (\text{III.67})$$

Based on Gaussian beam theory, the complex harmonic beam parameter q_n also verifies:

$$\frac{1}{q_n(z - z_n^r)} = \frac{1}{R_n(z - z_n^r)} - i \frac{\lambda_0}{n\pi w_n^2(z - z_n^r)} \quad (\text{III.68})$$

Based on equation III.63, the harmonic radius of curvature **at the emission plane** reads:

$$R(-z_{rn}) = f_p / \cos(\theta) \quad (\text{III.69})$$

The harmonic waist $w_n = w_n(0)$ and z_n^r can be deduced using equations III.64, III.68, III.67, III.66 and the relationship

$$\frac{1}{2R_n(-z_n^r)} = \frac{1}{2f_p / \cos(\theta)} - \frac{1}{2R_z}$$

Assuming a constant laser spatial phase (i.e. $\frac{1}{R_z} = 0$), we have:

$$\begin{aligned} z_n^r &= \frac{n\pi w_n^2}{\lambda_0} \frac{n\Psi}{1 + (n\Psi)^2} \\ w_n &= \frac{w_n(-z_n^r)}{\sqrt{1 + (n\Psi)^2}} \\ \Psi &= \frac{2\pi}{\cos(\theta)} \frac{w_n^2(z_n^r)}{w_0^2} \frac{2L \cos^2(\theta)}{\lambda_0} \end{aligned} \quad (\text{III.70})$$

From equations III.65 and III.70 one can show that:

$$\theta_n = \theta_n^0 \sqrt{1 + (n\Psi)^2} \quad (\text{III.71})$$

where θ_n^0 is the harmonic free divergence (i.e. the harmonic divergence in the absence of a PM curvature):

$$\theta_n^0 = \frac{\lambda_0}{n\pi w_n(-z_n^r)} \quad (\text{III.72})$$

In conclusion, when an ultra-intense laser pulse irradiates a solid target, a train of attosecond pulses of light, associated with Doppler harmonics in the frequency domain is generated by the relativistically

oscillating PM. One of the possible future applications of plasma mirrors resides in exploiting these ultra-fast light pulses in pump-probe like experiments. However, one of the main difficulties that still need to be addressed in order to develop a reliable source of X-UV attosecond light pulses from PMs is the difficulty of producing unique attosecond light pulses from those harmonics. One possible path towards this goal is the so-called attosecond lighthouse effect proposed a few years ago [26]. This scheme relies on laser spatio-temporal couplings to angularly separate attosecond light pulses in the far-field.

III.5 The attosecond lighthouse effect

The basic principle of the attosecond lighthouse effect is sketched in figure III.12. Panel (a) represents the typical setup used for a HHG experiment employing a standard Gaussian laser beam. In this case, the plasma mirror emits X-UV radiations in the form of a collimated train of attosecond light pulses along the specular direction. Panel (b) represents a HHG process with the so-called attosecond lighthouse effect. In this case, the incident laser field is distorted such that its wavefronts are continuously rotating in the incidence plane. When the laser-plasma interaction is taking place, a single attosecond light pulse is emitted at each laser optical cycle, whose propagation direction is given by the instantaneous laser wave vector (that has a propagation direction varying over time). If the laser wavefront rotation (WFR) velocity is high enough, then the emitted attosecond pulses are sufficiently angularly separated and a unique attosecond light pulse can be spatially filtered in the far-field by placing a slit along its path. This scheme is very general and is not restricted to HHG on plasma mirrors. For instance, it has already been applied to produce angularly separated attosecond light pulses from PM in the CWE regime [51] and on gaseous medias [52].




Figure III.12: Schematic representation of the lighthouse principle. In the presence of WFR, attosecond light pulses are emitted with slightly different directions. A unique attosecond light pulses can then be spatially filtered in the far-field

The attosecond lighthouse effect is based on a controlled distortion of the spatio-temporal profile of the laser pulse. This particular spatio-temporal coupling is called wavefront rotation. In this section, we formally introduce the notion of spatio-temporal couplings of light fields and show how to obtain,

in practice, wavefront rotation at laser focus. Finally, we review the necessary conditions to angularly split an attosecond pulses train into isolated attosecond pulses through the lighthouse effect.

III.5.1 Characterization of ultra-short laser pulses

We consider a laser electromagnetic wave, propagating along the z direction, with a wave vector \vec{k}_0 . This wave can be fully characterized, given its electric field spatio-temporal profile $E(x, z = z_0, t)$ at a fixed position $z = z_0$. The electric field can be characterized by its spatio-temporal phase $\varphi(x, t)$ and amplitude $A(x, t)$. Equivalently, this field admits three additional representations in different dual domains: (x, ω) , (k_x, t) , (k_x, ω) .

Space-frequency domain The space-frequency representation obtained by the Fourier transform of the spatio-temporal field over the time variable. Omitting the longitudinal position notation $z = z_0$, this representation can be expressed as:

$$\begin{aligned}\hat{E}(x, \omega) &= A(x, \omega)e^{i\varphi(x, \omega)} \\ &= \int_{-\infty}^{+\infty} E(x, z = z_0, t)e^{-i\omega t} dt\end{aligned}\quad (\text{III.73})$$

$A(x, \omega)$ denotes the spatio-spectral amplitude of the electromagnetic field while $\varphi(x, z = z_0, \omega)$ stands for its spatio-spectral phase. This representation gives an insight about the spectral distribution of field at different transverse positions x .

Spatial frequency-time domain The electromagnetic field may as well be expressed in the (k, ω) space. This representation is given by the Fourier transform of the spatio-temporal field over the transverse spatial variable.

$$\begin{aligned}\hat{E}(k, t) &= A(k, t)e^{i\varphi(k, t)} \\ &= \int_{-\infty}^{+\infty} E(x, t)e^{-ikx} dx\end{aligned}\quad (\text{III.74})$$

This representation is particularly interesting as it allows to express the electromagnetic field at the focal spot of an optics of focal length f based on Fraunhofer diffraction. For a narrow spectral bandwidth $\frac{\Delta k}{k_L} \ll 1$, the field after focusing reads:

$$E(x_f, t) \propto \hat{E}\left(\frac{k_L x_f}{f}, t\right)\quad (\text{III.75})$$

Spatial frequency-temporal frequency domain Finally, it is possible to express the electromagnetic fields in the (k, ω) space, given by the spatio-temporal Fourier transform of the field:

$$\begin{aligned}\hat{E}(k, \omega) &= A(k, \omega)e^{i\varphi(k, \omega)} \\ &= \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} E(x, t)e^{-i\omega t - ikx} dx dt\end{aligned}\quad (\text{III.76})$$

This representation informs on the angular spectral distribution with respect to the propagation direction z . In fact, the propagation angle θ can be expressed as $\theta(k_x, \omega) = \arcsin\left(\frac{ck_x}{\omega}\right)$. Thus, the (k, ω) representation informs on the angular intensity distribution of the laser pulse far from the laser waist:

$$I(\theta, \omega) \simeq \left| \hat{E}(k/k_L, \omega) \right|^2$$

III.5.2 First order spatio-temporal couplings in ultra short lasers

In section III.2, the laser spatial and temporal dependencies have been expressed separately. This does not always hold true and mutual dependencies between spatial and temporal laser characteristics actually arises in typical ultra-short laser systems. An ultra-short laser pulse is said to exhibit spatio-temporal couplings when:

$$E(x, z, t) \neq A_1(x) \exp(\varphi_1(x)) A_2(t) \exp(\varphi_2(t)) \quad (\text{III.77})$$

Spatio-temporal couplings can affect the laser amplitude and/or phase profiles.

$$\begin{aligned} \varphi(x, t) &\neq \varphi_1(x) \varphi_2(t) \\ A(x, t) &\neq A_1(x) A_2(t) \end{aligned} \quad (\text{III.78})$$

As the field admits in total four different representations in four different spaces, one could define 8 spatio-temporal couplings, each affecting the field phase or amplitude in one of these four spaces. As the different field representations are related via Fourier transforms, the underlying spatio-temporal couplings are interdependent. In this thesis, we restrict our study to Gaussian pulses exhibiting first order spatio-temporal couplings following the work of Aktürk *et al* [53]. In this case, the spatio-temporal field profile can be expressed as:

$$\begin{aligned} E(x, t) = E_0 \exp \left[- \left(\frac{1}{w_0^2} + i \frac{k_0}{2R} \right) x^2 \right] \times \\ \underbrace{\exp[(\xi + i\zeta)xt]}_{\text{first order coupling}} \times \\ \exp \left[- \left(\frac{1}{\tau^2} + i\beta \right) t^2 \right] \exp[i\omega_0 t] \end{aligned} \quad (\text{III.79})$$

with $\{w_0, R, \tau, \beta, \xi, \zeta\} \subset \mathbb{R}$. The spatio-temporal coupling is said to be of first order because the second order derivatives of the coupling term $(\xi + i\zeta)xt$ are equal to zero:

$$\partial_t^2 \left((\xi + i\zeta)xt \right) = \partial_x^2 \left((\xi + i\zeta)xt \right) = 0$$

The parameters involved in equation III.79 are defined as follows:

- β is called temporal chirp. It induces a varying instantaneous laser frequency.
- ξ and ζ are spatio-temporal coupling parameters, respectively affecting the amplitude and the phase of an ultra-short laser pulse.
 - ξ , is associated to the so-called **pulse-front tilt** (PFT) coupling (cf left panel figure III.13). It results in delayed arrival times of the laser amplitude peak across the transverse direction x . Consequently, the laser wave-fronts and pulse-fronts are tilted. Note that the pulse-front is defined as the variation of the time delay $t(x_0)$ of the maximum laser intensity for each transverse position x_0 .

$$\text{Pulse-Front} = \{(x, t), \forall x, t = \operatorname{argmax}_{t_0} (|E(x, t_0)|)\}$$

- ζ leads to a spatio-temporal coupling affecting the phase. In this case, the laser wavefronts rotate over time. This coupling is called **wave front rotation** (WFR) (cf right panel figure III.13). It is the basic mechanism responsible for the attosecond lighthouse effect.

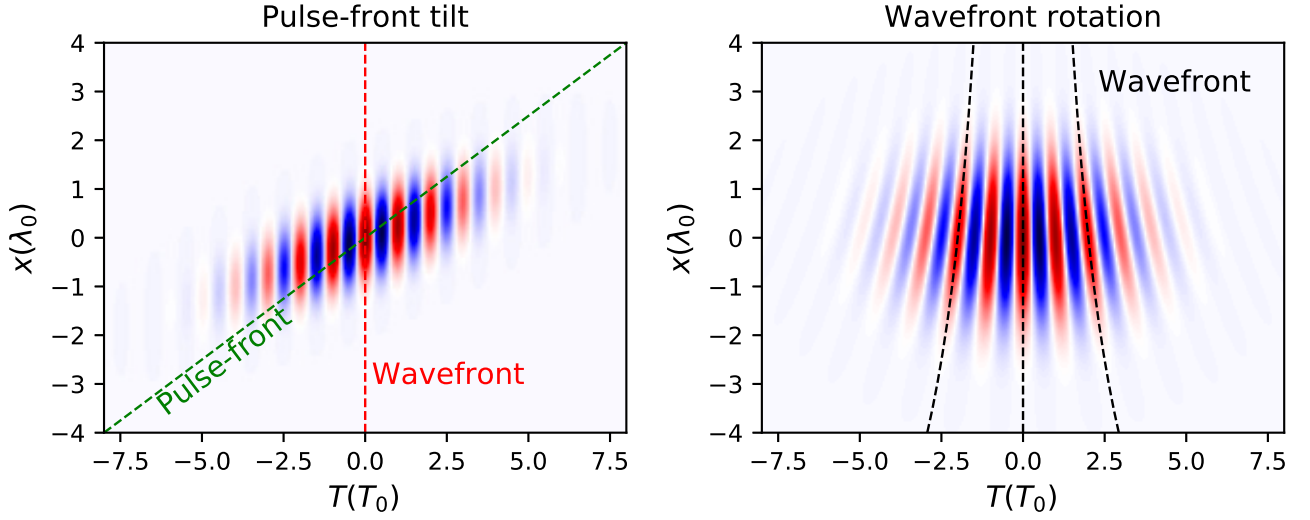


Figure III.13: Pulse-front tilt (left panel) and wavefront rotation (right panel) in the (x, t) space

Equation III.79 can be rewritten in a more general form:

$$E(x, t) = E_0 \exp \left[a_{xx}x^2 + a_{tt}t^2 + a_{xt}xt \right] \quad (\text{III.80})$$

with $\{a_{xx}, a_{tt}, a_{xt}\} \subset \mathbb{C}$ and:

- $\text{Re}(a_{xx}) = -\frac{1}{w_0^2} \Leftrightarrow$ beam waist.
- $\text{Im}(a_{xx}) = -\frac{ik}{2R} \Leftrightarrow$ wavefront curvature (WFC).
- $\text{Re}(a_{tt}) = -\frac{1}{\tau^2} \Leftrightarrow$ laser pulse duration.
- $\text{Im}(a_{tt}) = \beta \Leftrightarrow$ temporal chirp.
- $\text{Re}(a_{xt}) = \xi \Leftrightarrow$ pulse-front tilt.
- $\text{Im}(a_{xt}) = \zeta \Leftrightarrow$ wave-front rotation.

III.5.3 Different representations of spatio-temporal couplings

In the other representation spaces, the electric component of the laser field can be expressed with the same form as equation III.80:

$$\begin{aligned} \hat{E}(x, \omega) &\propto E_0 \exp \left[b_{xx}x^2 + b_{\omega\omega}\omega^2 + b_{x\omega}x\omega \right] \\ \hat{E}(k, \omega) &\propto E_0 \exp \left[c_{kk}k^2 + c_{\omega\omega}\omega^2 + c_{k\omega}k\omega \right] \\ \hat{E}(k, t) &\propto E_0 \exp \left[d_{kk}k^2 + d_{tt}t^2 + d_{kt}kt \right] \end{aligned} \quad (\text{III.81})$$

The spatio-temporal couplings are determined in each representation space by the coefficients $b_{x\omega}$, $c_{k\omega}$, d_{kt} . The next table III.1 summarizes the relationships between different coupling parameters:

	(x, t)	(x, ω)	(k, ω)	(k, t)
$a_{xt} =$	a_{xt}	$\frac{i}{2} \frac{b_{x\omega}}{b_{\omega x}}$	$\frac{1}{4} \frac{c_{k\omega}}{c_{kk}c_{\omega\omega} + c_{k\omega}^2}$	$-\frac{i}{2} \frac{d_{kt}}{d_{kk}}$
$b_{x\omega} =$	$\frac{i}{2} \frac{a_{xt}}{a_{tt}}$	$b_{x\omega}$	$\frac{i}{2} \frac{c_{k\omega}}{c_{kk}}$	$\frac{1}{4} \frac{d_{kt}}{d_{kk}d_{tt} + d_{kt}^2}$
$c_{k\omega} =$	$\frac{1}{4} \frac{a_{xt}}{a_{xx}a_{tt} + a_{xt}^2}$	$\frac{i}{2} \frac{b_{x\omega}}{b_{xx}}$	$c_{k\omega}$	$-\frac{i}{2} \frac{d_{kt}}{d_{tt}}$
$d_{kt} =$	$\frac{i}{2} \frac{a_{xt}}{a_{xx}}$	$\frac{1}{4} \frac{b_{x\omega}}{b_{xx}b_{\omega\omega} + b_{x\omega}^2}$	$\frac{i}{2} \frac{c_{k\omega}}{c_{\omega\omega}}$	d_{kt}

Table III.1: Relationship between spatio-temporal couplings in different domains

The relationships between other parameters (a_{xx} , a_{tt} , b_{xx} , $b_{\omega\omega}$, c_{kk} , $c_{\omega\omega}$, d_{kk} , d_{tt}) are given in appendix B.

Analyzing spatio-temporal couplings coefficients in different representation spaces brings additional insights to the physical origin of each coupling mechanism.

Let us consider a Gaussian beam at focus (i.e. with flat wavefronts $\text{Im}(a_{xx}) = 0$), free from temporal chirp ($\text{Im}(a_{tt}) = 0$). In this case, the WFR coupling ($\text{Im}(a_{xt}) \neq 0$) implies an amplitude coupling in the (x, ω) space ($\text{Re}(b_{x\omega}) \neq 0$). The central laser frequency ω varies as a function of the transverse position x . This is the definition of the spatial chirp (i.e. a space-frequency coupling). Moreover, WFR translates into an amplitude spectro-temporal coupling in the (k, t) space (i.e. $\text{Re}(d_{kt}) \neq 0$). This means that the laser wave vector is time-dependent. WFR is actually associated to the rotation of the laser wave vector over time: different laser wavefronts propagate along different directions.

On the other hand, a pulse-front tilt coupling in the spatio-temporal domain $\text{Re}(a_{xt}) \neq 0$ gives rise to an amplitude coupling in the (k, ω) space ($\text{Re}(c_{k\omega}) \neq 0$). As already mentioned, the (k, ω) space representation illustrates the angular spectral distribution of the laser field. A non null $\text{Re}(c_{k\omega})$ means that different frequencies constituting the laser wave packet propagate along different angular directions $\theta \simeq \frac{ck_x}{\omega}$. This spatio-temporal coupling is commonly called angular dispersion (AGD) [54].

Note also that angular dispersion ($\Leftrightarrow \text{Re}(c_{k\omega}) \neq 0$) could emerge even in the absence of pulse-front tilt ($\text{Re}(a_{xt}) = 0$). Indeed, as shown by table III.5.3, the presence of wavefront curvature ($\text{Im}(a_{xx}) \neq 0$) and wavefront rotation ($\text{Im}(a_{xt}) \neq 0$) may induce angular dispersion through the relationship:

$$c_{k\omega} = \frac{1}{4} \frac{a_{xt}}{a_{xx}a_{tt} + a_{xt}^2}$$

This case is illustrated by figure III.14. The left panel represents the spatio-temporal profile of a laser beam exhibiting WFR and wavefront curvature. The amplitude the spectrum of this field in the (k, ω) space is sketched in the left panel. The spatio-temporal couplings translate into a tilted spectrum amplitude in the (k, ω) space (i.e. an amplitude coupling in this space) associated to AGD.

It is important to mention that the wavefront curvature resulting from propagation in vacuum of a laser pulse, exhibiting wavefront rotation only (no initial PFT, no temporal chirp and no wavefront curvature) at focus **does not** induce angular dispersion. Indeed, assuming a short laser bandwidth, and under the paraxial approximation, free propagation in vacuum can be approximated by the Huygens-Fresnel method:

$$\hat{E}(k_x, \omega, z = z_0) = \hat{E}(k_x, \omega, z = 0) \exp\left(-iz_0 \frac{k_x^2}{2k_0}\right) \quad (\text{III.82})$$

where $k_0 = \frac{\omega_0}{c}$ is the laser wave vector. Hence, when the pulse propagates in vacuum, the coupling parameter in the (k, ω) space $c_{k\omega}$ is not modified. Therefore, the angular dispersion remains unchanged. In appendix A, we show how free propagation in vacuum simultaneously gives rise to PFT, wavefront curvature and temporal chirps such that the AGD is kept null.

However, inducing angular dispersion by curving the wavefronts of a rotating laser pulse is still possible if the wavefront curvature results from an external distortion of the laser wavefront. Such a distortion may be the result of the plasma mirror curvature in the case of Doppler HHG, or the laser intensity spatial variations in the case of CWE harmonics.




Figure III.14: Left panel: Spatio-temporal profile of an ultra-short laser pulse exhibiting WFR and wavefront curvature. The spectro-spectral representation of the same pulse. The tilt in the (k, ω) space is associated to AGD.

III.5.4 Practical realization of the attosecond lighthouse effect

The attosecond lighthouse effect relies on driving a HHG process using a laser pulse with WFR. As we detail now, inducing wavefront rotation into ultra-short laser pulses can be easily achieved in

CPA-based lasers.

The spectro-temporal representation in the (k, t) space of a laser pulse under the paraxial approximation can be interpreted as the spatio-temporal profile of the same pulse at the focal plane of a focusing optics. Therefore, spatio-temporal coupling parameter pertaining to the pulse at focus expressed in the (x, t) space is proportional to the spatio-temporal coupling before focusing when expressed in the (k, t) space (with a real proportionality coefficient) :

$$a_{xt}^{AF} \propto d_{kt}^{BF} \quad (\text{III.83})$$

where the acronyms AF and BF stand for at focus and before focusing respectively.

Consequently, if the initial pulse (before focusing) exhibits pulse-front tilt coupling ($\text{Re}(a_{xt}) \neq 0$, $\text{Im}(a_{xt}) = 0$), in the absence of wavefront curvature ($\text{Im}(a_{xx}) = 0$), we have $d_{kt}^{BF} \in i\mathbb{R}$. Hence, the spatio-temporal coupling parameter at focus a_{xt}^{AF} is pure imaginary, so the initial pulse-front tilt is converted into wave-front rotation at focus. Figure III.15 physically illustrates how a PFT may induce a WFR at focus [55]. Panel (a) represents the focusing of a laser pulse in the presence of PFT. Due to the spatio-temporal coupling in the initial pulse, different laser wavefronts across the laser pulse reach different spots of the focusing lens surface. Therefore, the laser wavefronts, not reaching the lens at its center are tilted at focus which results in the apparent rotation of the laser wavefronts. Equivalently, panel (b), shows how the coupling conversion can be explained by interpreting the PFT as angular dispersion. In this case, different colors constituting the initial pulse impinge the lens with different angles (due to AGD). Therefore, each wavelength is focused on a different transverse position of the lens focal plane. At focus, the resulting field exhibits a spatial chirp (frequency dependence to the transverse position) which is equivalent to a wavefront rotation. This behavior suggests that WFR can be triggered by focusing a laser pulse with PFT. So the question is, how to induce and control the pulse-front tilt before the focus?

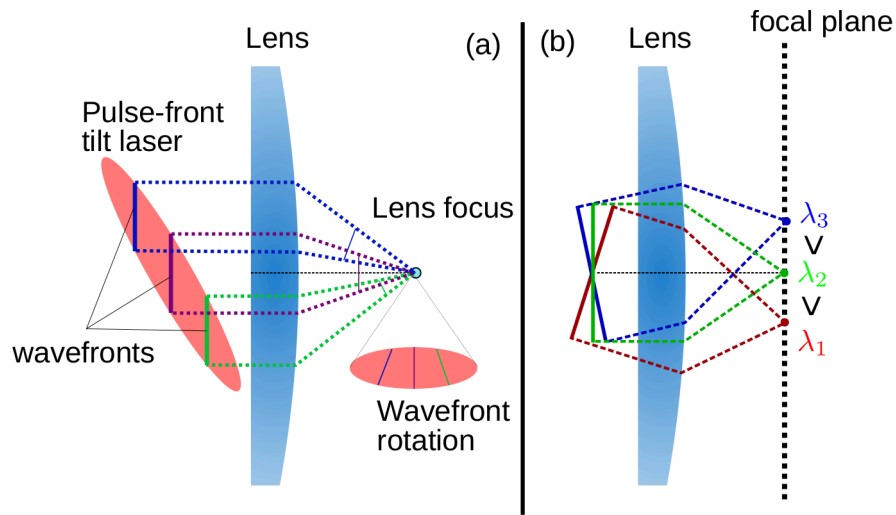


Figure III.15: Focusing a laser pulse in the presence of PFT/AGD. Panel (a), in the presence of PFT, the laser wavefronts are tilted with respect to the pulse-front. At the focus of the lens, the initially parallel wavefronts become tilted with respect to each other, giving rise to WFR. Panel (b): in the presence of AGD, each wavelength constituting the laser wave packet impinges the lens surface with a different angle. Hence, each frequency is focused on a different position of the lens focal plane, giving rise to angular dispersion.

III.5.5 Inducing pulse-front tilt

The real challenge to control laser wavefront rotation velocity mostly resides in inducing and manipulating the amount of pulse-front tilt before focusing. Figure III.16 schematically illustrates how an ultra-short regular laser pulse refraction inside a dispersive optical prism gives rise to PFT. Inside the prism, the laser pulse-front travels at the group velocity, while the laser wavefront travels at the phase velocity. Due to the dispersive nature of the prism, the group and phase velocities of light inside the prism are not equal. Therefore, pulse-front (which travels at the group velocity) and the wavefront (which travels at the phase velocity) are tilted with respect to each other by the prism.

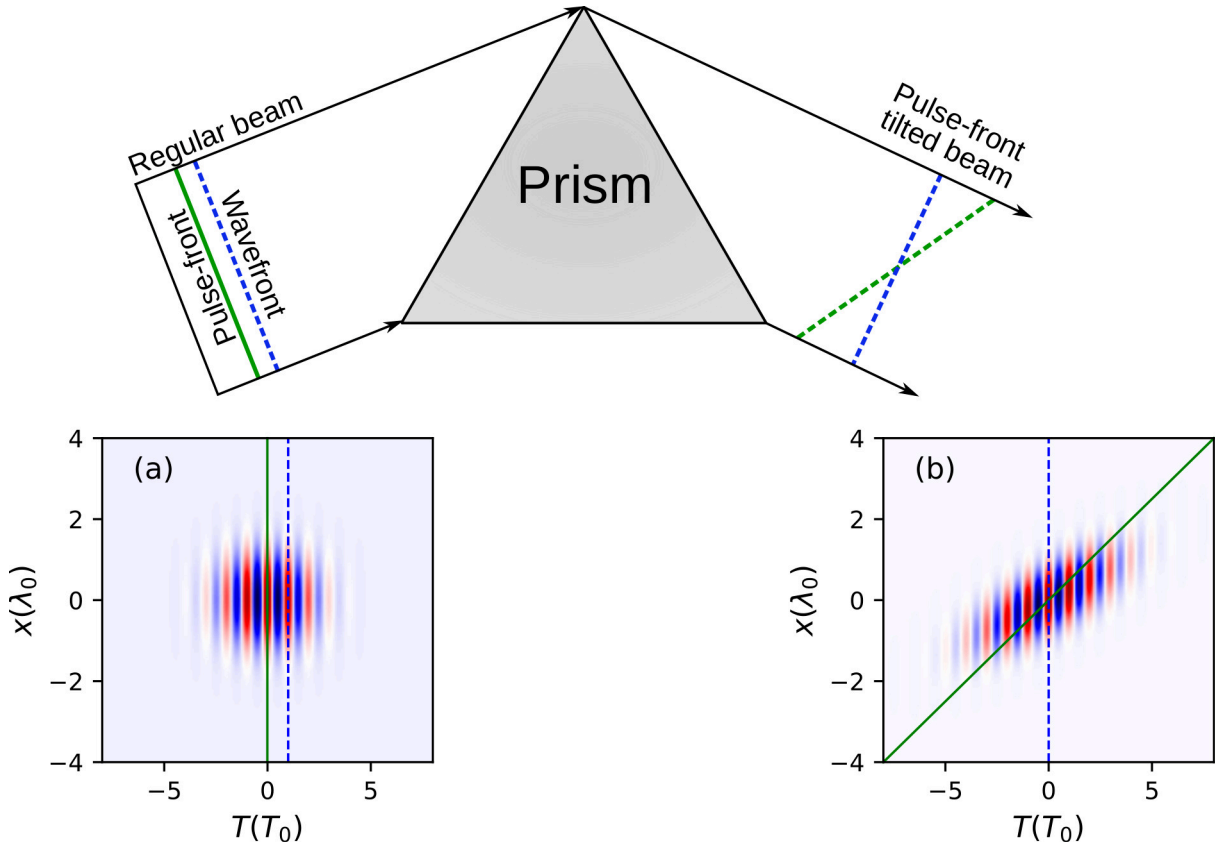


Figure III.16: Pulse-front tilt induced by a dispersive prism.

In reference [54], the author derived a general relationship between the angular dispersion and the pulse-front tilt induced by the propagation of an ultra-short laser pulse in any dispersive medium.

This relationship reads:

$$\begin{aligned} \tan(c\xi) &= -\lambda_0 \frac{d\Gamma(\lambda)}{d\lambda} \\ \xi &= \arctan\left(\lambda_0 \frac{d\Gamma(\lambda)}{d\lambda}\right)/c \end{aligned} \quad (\text{III.84})$$

where $\Gamma(\lambda)$ is the propagation angle gap between the spectral component of wavelength λ and the central laser spectral component (with wavelength λ_0) of the resulting pulse. The quantity $\frac{d\Gamma(\lambda)}{d\lambda}$ characterizes the angular dispersion effect of the dispersive media.

In practice, introducing a pulse-front tilt into an ultra-intense ultra-short laser pulse relies on misaligning the last grating of a laser system. This grating is usually used to recombine different spectral components of the pulse in typical CPA-based chains [1]. Figure III.17 schematically illustrates the light diffraction in such a setup.




Figure III.17: Diffraction on a misaligned double-grating.

Equation III.84 suggests that the pulse-front tilt induced by this setup can be deduced from its angular dispersion. We assume that the two gratings have the same step denoted a . We only consider first order diffractions on both gratings. The diffraction grating formulas read:

$$\begin{aligned}\sin(i) + \sin(r) &= \frac{\lambda}{a} \\ \sin(i - \Theta) + \sin(\Gamma) &= \frac{\lambda}{a}\end{aligned}\quad (\text{III.85})$$

Differentiating both equations with respect to λ gives:

$$\begin{aligned}\frac{dr}{d\lambda} &= \frac{1}{a \cos(r)} \\ \frac{d\Gamma}{d\lambda} &= \frac{1}{a} - \cos(r - \Theta) \frac{dr}{d\lambda}\end{aligned}\quad (\text{III.86})$$

after few manipulations we get:

$$\frac{d\Gamma}{d\lambda} = \frac{1 - \cos(r - \Theta)/\cos(r)}{a \sqrt{1 - (\lambda/a - \sin(r - \Theta))^2}} \quad (\text{III.87})$$

under the paraxial approximation, and using the relationship of equation III.84, we get:

$$\tan(c\xi) = -\frac{\lambda_0 \tan(r)}{a \cos(i)} \Theta \quad (\text{III.88})$$

Concretely, pulse-front tilt can be controlled at will by slightly misaligning a double-gratings in a CPA laser chain via the relationship III.88.

III.5.6 From PFT to WFR

In this section, we characterize a Gaussian laser with wavefront rotation resulting from the focusing of an ultra-short laser with PFT. Let f be the focal distance of the focusing optics.

Before focus, the pulse-front tilted electric field expression reads:

$$E(x_i, t) = E_0 \exp\left(-\left(\frac{t - \xi x_i}{\tau_i}\right)^2 - \left(\frac{x_i}{w_i}\right)^2 - i\omega_0 t\right) \quad (\text{III.89})$$

in the (k, t) space, this expression becomes:

$$\begin{aligned}\hat{E}(k_x, t) &= \int_{-\infty}^{\infty} E(x_i, t) \exp(-ik_x x_i) dx_i \\ \hat{E}(k_x, t) &= E^0 \exp\left(-w_i^2 \frac{k_x^2}{1 + (w_i \xi / \tau_i)^2} - \frac{t^2}{\tau_i^2 (1 + (w_i \xi / \tau_i)^2)}\right) \times \\ &\quad \exp\left(i\omega t + ik_x t \frac{\xi w_i^2}{\tau_i^2 (1 + (w_i \xi / \tau_i)^2)}\right)\end{aligned}\quad (\text{III.90})$$

The field expression at focus can be computed by the Fraunhofer diffraction:

$$E^{AF}(x_f, t) \propto \hat{E}\left(k_x = \frac{x_f k_0}{f}, t\right) \quad (\text{III.91})$$

using III.90, we get:

$$\begin{aligned}E(x_f, t) &= E_0^f \exp\left(-\frac{x_f^2}{w_f^2 (1 + (\xi w_i / \tau_i)^2)}\right) \times \\ &\quad \exp\left(-\frac{t^2}{\tau_f^2 (1 + (\xi w_i / \tau_i)^2)}\right) \times \\ &\quad \exp(i\omega_0 t + i\zeta t x)\end{aligned}\quad (\text{III.92})$$

with

$$w_f = w_0 \sqrt{1 + (\xi w_i / \tau_i)^2}$$

and

$$w_0 = \frac{\lambda_0 f}{\pi w_i}$$

w_f is the effective beam waist at focus while w_0 is the usual laser waist at focus in the absence of PFT. Moreover, we have:

$$\tau_f = \tau_i \sqrt{1 + (\xi w_i / \tau_i)^2}$$

the laser temporal duration at focus and:

$$\zeta = \frac{2\xi w_i}{w_0 \tau_i^2 (1 + (\xi w_i / \tau_i)^2)}$$

is the laser wavefront rotation parameter.

The PFT before focus increases the pulse duration and waist at focus by a factor of $\sqrt{1 + (\xi w_i / \tau_i)^2}$. Assuming that the laser energy is conserved, the laser amplitude at focus $|E_0^f|$ reads:

$$\frac{|E_0^f|}{|E^0|} = \sqrt{\frac{w_i}{w_0}} \frac{1}{\sqrt{1 + (\xi w_i / \tau_i)^2}}$$

We now compute the laser wavefront rotation velocity at focus. We first define the instantaneous laser direction of propagation as $\beta \simeq \frac{k_{\perp}}{k_0}$ with $k_{\perp}(t) = k_x \frac{d(\zeta t x)}{dx}$. Finally, we define the laser wavefront rotation velocity V_r as the temporal derivative of its instantaneous direction of propagation.

$$\begin{aligned}V_r(\xi) &= \frac{d\beta}{dt} \\ V_r(\xi) &= \frac{\zeta}{k_0} \\ V_r(\xi) &= \frac{w_i^2 \xi}{f \tau_i^2 (1 + (\xi w_i / \tau_i)^2)}\end{aligned}\quad (\text{III.93})$$

Figure III.18, shows the evolution of the laser WFR velocity as a function of the PFT parameter before focusing ξ . The wavefront rotation reaches a peak at $\xi = \xi_{max} = \frac{\tau_i}{w_i}$ before slowly decreasing. The maximum WFR velocity reads:

$$\begin{aligned}
 V_r^{max} &= \frac{w_i}{2\tau_i f} \\
 \tau_f &= \sqrt{2}\tau_i \\
 w_f &= \sqrt{2}w_0 \\
 \frac{|E_0^f|}{|E^0|} &= \frac{1}{\sqrt{2}}\sqrt{\frac{w_i}{w_0}}
 \end{aligned} \tag{III.94}$$




Figure III.18: Wavefront rotation velocity as a function of the pulse-front tilt before focus.

Maximizing the laser wavefront rotation velocity at focus induces a temporal and spatial beam sizes increase by a factor of $\sqrt{2}$ and a laser intensity decrease by 50%. Note also that the wavefront rotation velocity is inversely proportional to the laser temporal duration before focus τ_i . Therefore, longer laser pulses would sustain lower maximum wavefront velocities.

III.5.7 Separation condition of attosecond light pulses

Figure III.19 illustrates under which conditions the use of a rotating laser pulse leads to angular separation of attosecond light pulses in the far-field.

Assuming that the emission direction of attosecond light pulses follows the instantaneous driving laser wave-vector direction, the angular separation between two successive attosecond light pulses is $\Delta\theta = \Delta T V_r$ where ΔT is the delay between the emission times of two successive attosecond light pulses. In the context of Doppler HHG, a single attosecond light pulse is emitted at each laser optical cycle. Therefore we have

$$\Delta\theta = T_0 V_r$$

As the train of attosecond light pulses diffracts in vacuum, spatially filtering a single attosecond pulse in the far-field with a slit is only possible when $\Delta\theta$ is larger than the angular spread of individual attosecond pulses. Otherwise, the filtered signal will be made of two (or more) undesired parasite spikes coming from attosecond pulses close to the targeted pulse before the slit. The higher the temporal contrast of the filtered signal, the better the angular separation is. Additionally, the smaller the slit is, the higher is the contrast between the central pulse and its satellite pulses. However, this also lowers the total energy in the central pulse. Assuming that harmonic beam divergences constitute a good measurement of their angular spread, we consider that a good separation is achieved when:

$$\frac{\Delta\theta}{\theta_n} \gtrsim 1 \quad (\text{III.95})$$

where θ_n is the divergence of a single attosecond pulse. This separation criterion is arbitrary, and depending on the potential applications and use of attosecond light pulses, a higher separation ratio $\frac{\Delta\theta}{\theta_n}$ might be required to increase the temporal contrast of the filtered signal.




Figure III.19: Attosecond pulses separation with the lighthouse scheme. The separation is only possible when the angular divergence θ_n of a single attosecond light pulse is smaller than the angular separation between successive pulses $\Delta\theta$.

Successful experimental demonstrations of attosecond lighthouse effect, whether on gaseous media or on plasma mirrors in the non-relativistic regime were performed using extremely short laser pulses (FWHM = $5fs$) in order to induce a high laser WFR velocity. Nevertheless, this scheme has never been transposed to the relativistic regime to produce isolated attosecond light pulses. The main difficulty arising in the context of Doppler HHG results from the important divergence of high order harmonics in this regime. Therefore, the required WFR velocity is only achievable when using extremely short laser pulses (of one or two cycles only). Unfortunately, amplifying laser pulses with such broad bandwidths to reach relativistic intensities remains a hard challenge to solve for current laser technologies.

Chapter IV

Basics of the PIC algorithm

IV.1 Numerical modelling in plasma physics	48
IV.2 Maxwell-Vlasov equations	48
IV.3 The Particle-In-Cell (PIC) method	49
IV.3.1 WARP + PICSAR PIC framework	50
IV.3.2 Discretization of PIC quantities	51
IV.3.3 Current/charge deposition	52
IV.3.4 Fields interpolation	53
IV.3.5 Particles pusher	54
IV.3.6 Maxwell Solver	54
IV.4 FFT-based Pseudo-spectral Maxwell solvers	60
IV.4.1 Pseudo-Spectral Analytical Time Domain	60
IV.4.2 Parallelization of the pseudo-spectral methods	63
IV.4.3 Truncation error and ultra-high order solvers	63
IV.5 Absorbing boundary conditions with PSATD	66

IV.1 Numerical modelling in plasma physics

The plasma state describes a broad range of physical systems with wide variety of parameters (temperature, density, pressure, ionization degree ...), thus various theoretical descriptions of plasma states are possible, depending on the physical mechanisms at play.

In this context, we can identify three possible modelling approaches to describe plasma systems:

- Microscopic description: where each particle constituting the plasma medium is modelled individually. This modelling is equivalent to an N-body problem where each particle dynamics is governed by external forces induced by all other particles.
- Statistical (or kinetic) description: where plasma is described via the density distribution function of each of its species in the phase space (t, \vec{r}, \vec{p}) . In this case, a statistical description of the plasma medium is employed. Depending on the importance of binary collisions, this approach gives rise to the Boltzmann equation (for collisional plasmas) or Vlasov equation (for collision-less plasmas).
- Fluid description: when the plasma is near equilibrium, the plasma can be described via its density, mean velocity, temperature and pressure following a fluid approximation.

In the context of high order harmonic generation on solid targets, kinetic effects are dominant as plasmas are out of equilibrium. However, collisional effects can be neglected as the electrons collisions period is much larger than the total laser-plasma interaction duration. Therefore, the Vlasov equations, coupled with the Maxwell's equations, is a satisfying approach to describe both the electromagnetic fields and the plasma dynamics.

IV.2 Maxwell-Vlasov equations

The Maxwell-Vlasov equations are a PDE system that describe the evolution of the density distribution functions of each plasma species in the 6D phase space (x, y, z, p_x, p_y, p_z) in the collision-less regime. Let us consider a species s of mass m_s and charge q_s associated with a density distribution function f_s . It can be shown that under the collision-less hypothesis, the density distribution function is solution of the Vlasov equation [29]:

$$\frac{\partial f_s}{\partial t} + \frac{d\vec{r}}{dt} \cdot \nabla_{\vec{r}} f_s + \frac{d\vec{p}}{dt} \cdot \nabla_{\vec{p}} f_s = 0 \quad (\text{IV.1})$$

Assuming that the plasma is only prone to electromagnetic forces (neglecting gravity), the Vlasov equation is simply an advection equation whose characteristic lines (\vec{r}, \vec{p}) are solutions to the following equations:

$$\begin{aligned} \frac{d\vec{p}}{dt} &= q_s (\vec{E} + \vec{v} \wedge \vec{B}) \\ \frac{d\vec{r}}{dt} &= \vec{v} = \frac{\vec{p}}{\sqrt{m_s^2 + |\vec{p}|^2/c^2}} \end{aligned} \quad (\text{IV.2})$$

On the other hand, the electromagnetic fields are governed by the Maxwell's equations that we recall here:

$$\vec{\nabla} \cdot \vec{E} = \frac{\rho}{\varepsilon_0} \qquad \vec{\nabla} \wedge \vec{E} = -\frac{d\vec{B}}{dt} \qquad (\text{IV.3})$$

$$\vec{\nabla} \cdot \vec{B} = 0 \qquad \vec{\nabla} \wedge \vec{B} = \mu_0 \vec{J} + \frac{d\vec{E}}{c^2 dt} \qquad (\text{IV.4})$$

Maxwell's equations are coupled to the density distribution functions $f_s(\vec{r}, \vec{p}, t)$ via the source terms ρ and \vec{J} , which are related to f_s via the following relationship:

$$\begin{aligned} \rho &= \sum_s q_s \int f_s dp_x dp_y dp_z \\ \vec{J} &= \sum_s q_s \int f_s \frac{\vec{p}}{m_s \sqrt{1 + \frac{|\vec{p}|^2}{(m_s c)^2}}} dp_x dp_y dp_z \end{aligned} \qquad (\text{IV.5})$$

A Eulerian discretization of the Maxwell-Vlasov system given by equations [IV.1](#) [IV.4](#) and [IV.5](#) mandates to discretize 6 computation axes (3 space dimensions and 3 momentum dimensions) and would necessitate thousands of billions of cells even with a coarse grid sampling (say one hundred grid point per axis). The use of the Particle-In-Cell method represents a good numerical strategy to substantially reduce the computational complexity of the Maxwell-Vlasov system while still keeping an accurate kinetic description of the physics at play. In this case, only the spatial domain is discretized, while the density distribution function variation with respect to momentums is inferred via a Lagrangian formalism.

IV.3 The Particle-In-Cell (PIC) method

The PIC method [[56](#), [57](#)] is a particle-mesh approach to solve the Maxwell-Vlasov system. On the one hand, the electromagnetic quantities, obeying to the Maxwell's equations are discretized using a Eulerian description on a fixed simulation grid.

On the other hand, the density distribution function is computed using a Lagrangian approach. In this case, the distribution function is divided into N elementary distribution functions in the phase space (\vec{r}, \vec{p}) :

$$f_s(\vec{r}, \vec{p}, t) = \sum_{i=1}^N f_s^i(\vec{r}, \vec{p}, t)$$

These elementary functions, commonly called macro-particles, are supposed to have a finite extension in space with a definite momentum \vec{p} at all times:

$$f_s^i(\vec{r}, \vec{p}, t) = h(\vec{r} - \vec{r}_i(t)) \delta(\vec{p} - \vec{p}_i(t))$$

where h is a bounded-support function centered around 0, and $(\vec{r}_i(t), \vec{p}_i(t))$ are the coordinates in the phase space of the center of the i^{th} macro-particle at time t . These macro-particles can be seen as solid bodies driven by electromagnetic forces. Each macro-particle i , $(\vec{r}_i(t), \vec{p}_i(t))$ follows a well determined Vlasov characteristic line in the phase space, specified by equation [IV.2](#) and the initial conditions $(\vec{r}_i(t=0), \vec{p}_i(t=0))$. Therefore, the macro-particles dynamics is simply governed by the

relativistic momentum conservation principle. For this reason, macro-particles are seen as a collection of physical particles which are close to each other in the phase space (both in momentum and position).

After initializing the plasma macro-particles and initial fields in the simulation domain, the PIC algorithm is essentially composed of four main steps that are summarized here:

1. Maxwell solver: advance electromagnetic fields over one time step on the simulation grid, taking into account the electric current and the charge densities.
2. Fields interpolation: interpolate electromagnetic fields on each particle position.
3. Particle pusher: use the interpolated fields to update the particles momentums and positions.
4. Current/charge deposition: use the updated particles velocities and positions to compute the new currents and charge densities on the simulation grid.

The PIC algorithm is a simple, yet powerful numerical approach to tackle a large variety of plasma physics problems. Note that this representation is not an exhaustive description of all the capabilities of modern PIC codes which often include advanced physical modules for modelling additional physical phenomena such as ionization processes, quantum electrodynamics effects, and collisions, but these effects are out of the scope of this work and will not be further discussed.

A wide variety of numerical techniques can be used in order to perform each of the four PIC steps. In practice, PIC simulations are very computationally expensive as they require billions of macro-particles and grid cells for a full 3D geometry. In the following, we present the main numerical methods and parallelization strategies used in the PIC framework WARP+PICSAR to which I contributed during this thesis.

IV.3.1 WARP + PICSAR PIC framework

WARP+PICSAR software is a PIC code composed of two independent softwares, that are WARP¹ and PICSAR². WARP is a PIC code written in Fortran and Python languages and developed at the Lawrence Berkeley National Laboratory (LBNL), while PICSAR is a full Fortran library, co-developed by CEA and LBNL, it includes highly optimized PIC computing kernels and supports multiple levels of parallelisms. PICSAR is intended to be used as a PIC toolbox for other PIC codes, but it can also work as a standalone simulation package (although with limited input/output features). Currently, PICSAR can be used by SMILEI code [58] (to access to FFT based Maxwell solvers that will be presented later) and WARPX [59] to boost specific computations.

When using the WARP+PICSAR framework, all the computationally demanding operations are performed with the optimized PICSAR kernels, while the WARP code mainly serves to set the initial simulation setup, carry out the numerical diagnostics, and schedule the PIC computations. It also offers a user-friendly interface to design numerical experiments.

¹<https://blast.lbl.gov/blast-codes-warp/>

²<https://picsar.net/>

IV.3.2 Discretization of PIC quantities

The standard spatial/temporal discretization for field quantities (current/charge and electromagnetic components) in most PIC codes relies on the Yee lattice [60]. This scheme defines two spatially-staggered discretizations along each simulation axis. Each mesh grid quantity is discretized within a specific grid along each direction. This setup allows for a spatio-temporal leap-frog integration of Maxwell's equations. Figure IV.1 represents the fields discretization over the Yee lattice as defined in the WARP+PICSAR framework. Note that in the original Yee lattice, electric and magnetic fields are defined on half-time step shifted times. This consideration is adjusted in the WARP+PICSAR PIC code, depending on the Maxwell solver algorithm that is used (more details on that in the Maxwell solvers section).

On the other hand, the discretization of macro-particles momentum/positions usually used in most PIC codes relies on temporal staggering between position and momentum in order to enable a leap-frog integration of macro-particles dynamics as well (cf table IV.1).




Figure IV.1: Schematic representation of fields staggering on the Yee lattice.

Fields	x	y	z	t
E_x	$(i + \frac{1}{2})\Delta x$	$j\Delta y$	$k\Delta z$	$n\Delta t$
E_y	$i\Delta x$	$(j + \frac{1}{2})\Delta y$	$k\Delta z$	$n\Delta t$
E_z	$i\Delta x$	$j\Delta y$	$(k + \frac{1}{2})\Delta z$	$n\Delta t$
B_x	$i\Delta x$	$(j + \frac{1}{2})\Delta y$	$(k + \frac{1}{2})\Delta z$	$(n + \frac{1}{2})\Delta t$
B_y	$(i + \frac{1}{2})\Delta x$	$j\Delta y$	$(k + \frac{1}{2})\Delta z$	$(n + \frac{1}{2})\Delta t$
B_z	$(i + \frac{1}{2})\Delta x$	$(j + \frac{1}{2})\Delta y$	$k\Delta z$	$(n + \frac{1}{2})\Delta t$
J_x	$(i + \frac{1}{2})\Delta x$	$j\Delta y$	$k\Delta z$	$(n + \frac{1}{2})\Delta t$
J_y	$i\Delta x$	$(j + \frac{1}{2})\Delta y$	$k\Delta z$	$(n + \frac{1}{2})\Delta t$
J_z	$i\Delta x$	$j\Delta y$	$(k + \frac{1}{2})\Delta z$	$(n + \frac{1}{2})\Delta t$
ρ	$i\Delta x$	$j\Delta y$	$k\Delta z$	$(n + 1)\Delta t$
Particles	-	-	-	t
Position	-	-	-	$(n + 1)\Delta t$
Momentum	-	-	-	$(n + \frac{1}{2})\Delta t$

Table IV.1: Spatial and/or temporal positions of field quantities and macro-particles at the end of the n^{th} PIC iteration.

IV.3.3 Current/charge deposition

The current/charge deposition consists in retrieving the electromagnetic sources (ρ, \vec{J}) from the macro-particles distribution as shown by equation IV.5. Numerically, macro-particles are provided finite but non-null spatial extent to interact with the grid. In a discretized approach, numerically computing the charge density consists in computing $\rho_{i,j,k}^n$ in the following way:

$$\rho_{i,j,k}^n = \frac{1}{\Delta x \Delta y \Delta z} \sum_s q_s \sum_{l=1}^{N_s} W_l^s \cdot S(x_{(l,s)}^n, y_{(l,s)}^n, z_{(l,s)}^n)_{(i,j,k)} \quad (\text{IV.6})$$

with:

- $(\Delta x, \Delta y, \Delta z)$ are the spatial cell sizes along x,y,z directions.
- N_s the number of macro-particles of the s^{th} species.
- W_l is the weight of the l^{th} macro-particle of the s^{th} species. This quantity is equivalent to the number of real particles that the macro-particles represents.
- $(x_{l,s}^n, y_{l,s}^n, z_{l,s}^n)$ denote the spatial coordinates of one macro-particle at time step n.
- $S(i, j, k)$ is the spatial shape function centered on the node i, j, k . The shape function reflects the extent of macro-particles in space.

Depending on the spatial extent of shape functions $S(i, j, k)$, each macro-particles inducts a charge density on an arbitrary number of grid points. The shape functions used in the WARP+PIC SAR PIC code are those derived in [57], which guarantee a smooth evolution of the density distribution as the macro-particles move on the grid, and a bounded support contribution of each macro-particle to the total charge density. These shape functions are spline functions of arbitrary orders k , which are recursively defined as follows:

$$S^k(x) = S^1 \otimes S^{k-1}(x) \quad (\text{IV.7})$$

where S^1 is the top-hat function defined as:

$$S^1(x) = \begin{cases} 1 & \text{if } -\frac{\Delta x}{2} < x \leq \frac{\Delta x}{2} \\ 0 & \text{otherwise} \end{cases} \quad (\text{IV.8})$$

The evolution of shape functions for $k = 1..3$ in 1D are depicted in figure IV.2. The number of grid points involved in the charge deposition is equal to spline order k along each dimension, and its integral over space is equal to one cell size.




Figure IV.2: Shape functions for $k = 1$, $k = 2$, and $k = 3$.

The choice of the shape function order is a trade-off between the tolerated numerical noise and the computational cost: higher orders guarantee a smoother charge density deposition while requiring more computations.

The same computational method described to compute ρ can be used to compute J :

$$J_{x_{i+\frac{1}{2},j,k}}^{n+\frac{1}{2}} = \frac{1}{\Delta x \Delta y \Delta z} \sum_s q_s \sum_{l=1}^{N_s} W_l^s v_{l,s}^x \cdot S(x_{(l,s)}^n, y_{(l,s)}^n, z_{(l,s)}^n)_{(i+\frac{1}{2},j,k)} \quad (\text{IV.9})$$

However, this method does not guarantee that the continuity equation is verified:

$$\frac{\partial \rho}{\partial t} + \vec{\nabla} \cdot \vec{J} = 0 \quad (\text{IV.10})$$

A modified version of the current deposition using arbitrary shape function is employed. This method is known as the Esirkepov current deposition algorithm [61], it verifies the discretized continuity equation, considering second order derivatives spatial and temporal derivative operators:

$$\frac{\rho_{i,j,k}^{n+1} - \rho_{i,j,k}^n}{\Delta t} + \frac{J_{x_{i+\frac{1}{2},j,k}}^{n+\frac{1}{2}} - J_{x_{i-\frac{1}{2},j,k}}^{n+\frac{1}{2}}}{\Delta x} + \frac{J_{y_{i,j+\frac{1}{2},k}}^{n+\frac{1}{2}} - J_{y_{i,j-\frac{1}{2},k}}^{n+\frac{1}{2}}}{\Delta y} + \frac{J_{z_{i,j,k+\frac{1}{2}}}^{n+\frac{1}{2}} - J_{z_{i,j,k-\frac{1}{2}}}^{n+\frac{1}{2}}}{\Delta z} = 0 \quad (\text{IV.11})$$

IV.3.4 Fields interpolation

To solve the macro-particles dynamics, the Lorentz forces acting on the macro-particles are required. For this, the electromagnetic fields, initially defined on the mesh grid are interpolated on

the macro-particles positions employing shape functions, in the same fashion as the current/charge deposition step:

$$E_x^n(x_l, y_l, z_l) = \sum_{i,j,k} S((i + \frac{1}{2})\Delta x, j\Delta y, k\Delta z)_{(x_l, y_l, z_l)} E_x^n((i + \frac{1}{2})\Delta x, j\Delta y, k\Delta z) \vec{e}_x \quad (\text{IV.12})$$

Note that the shape functions used for both the fields interpolation and the current/charge deposition need to be the same to avoid unphysical self-exerted forces.[57].

IV.3.5 Particles pusher

Once both the electric and magnetic fields are known on each macro-particle, the macro-particles dynamics is updated by solving equation IV.2. Several numerical approaches are possible to integrate this equation of motion. The most commonly employed scheme in PIC codes is the well-known Boris pusher [56]. The Boris pusher updates the macro-particles momentums $p^{n-\frac{1}{2}}$ and positions x^n through a leap-frog integration:

$$\frac{p^{n+\frac{1}{2}} - p^{n-\frac{1}{2}}}{\Delta t} = q_s \left[E^n + \left(\frac{p^{n+\frac{1}{2}} + p^{n-\frac{1}{2}}}{2\gamma^n m_s} \right) \wedge B^n \right] \quad (\text{IV.13})$$

$$\gamma^n = \sqrt{1 + \left(\frac{p^{n-\frac{1}{2}} + p^{n+\frac{1}{2}}}{2m_s c} \right)^2}$$

$$\frac{x^{n+1} - x^n}{\Delta t} = \frac{p^{n+\frac{1}{2}}}{\sqrt{m_s^2 + \|p^{n+\frac{1}{2}}\|^2/c^2}} \quad (\text{IV.14})$$

The right hand side of equation IV.13 involves the unknown momentum $p^{n-\frac{1}{2}}$ and would therefore require an implicit integration. However, Boris showed that this integration can be performed explicitly by splitting the Lorentz force as follows:

$$\vec{F}_L = \vec{F}_1 + \vec{F}_2 + \vec{F}_3 = \underbrace{\frac{1}{2}\vec{F}_E}_{\frac{1}{2}\text{Electric force}} + \underbrace{\vec{F}_B}_{\text{Magnetic force}} + \underbrace{\frac{1}{2}\vec{F}_E}_{\frac{1}{2}\text{Electric force}}$$

and performing a three-steps integration of the equation of motion, using $\vec{F}_1, \vec{F}_2, \vec{F}_3$ successively.

IV.3.6 Maxwell Solver

While most modern Particle-In-Cell softwares rely on finite-different schemes to solve the Maxwell's equations, WARP+PICSAR offers the possibility to use more sophisticated, FFT-based, pseudo-spectral solvers along with original efficient parallelization strategies. First, we present the Finite-Different Time Domain solver employed by many PIC codes and discuss its limitations in the context to HHG simulations. Then, we present the pseudo-spectral solver, implemented in WARP+PICSAR code, along with the innovative parallelization approach initially proposed by J.L.Vay *et al* [25].

Finite-Difference Time Domain

The Finite-Difference Time Domain (FDTD) technique [60, 62] is based on a finite-difference scheme to compute numerical differentiations in both time and space using a low order stencil. The FDTD integration of Maxwell's equations reads as follows:

$$\begin{aligned}
 \frac{1}{c^2}(D_t E_{x_{i+\frac{1}{2},j,k}})^{n+\frac{1}{2}} &= (D_y B_z^{n+\frac{1}{2}})_{i+\frac{1}{2},j,k} - (D_z B_y^{n+\frac{1}{2}})_{i+\frac{1}{2},j,k} - \mu_0 J_{x_{i+\frac{1}{2},j,k}}^{n+\frac{1}{2}} \\
 \frac{1}{c^2}(D_t E_{y_{i,j+\frac{1}{2},k}})^{n+\frac{1}{2}} &= -(D_x B_z^{n+\frac{1}{2}})_{i,j+\frac{1}{2},k} + (D_z B_x^{n+\frac{1}{2}})_{i,j+\frac{1}{2},k} - \mu_0 J_{y_{i,j+\frac{1}{2},k}}^{n+\frac{1}{2}} \\
 \frac{1}{c^2}(D_t E_{z_{i,j,k+\frac{1}{2}}})^{n+\frac{1}{2}} &= (D_x B_y^{n+\frac{1}{2}})_{i,j,k+\frac{1}{2}} - (D_y B_x^{n+\frac{1}{2}})_{i,j,k+\frac{1}{2}} - \mu_0 J_{z_{i,j,k+\frac{1}{2}}}^{n+\frac{1}{2}} \\
 (D_t B_{x_{i,j+\frac{1}{2},k+\frac{1}{2}}})^{n+1} &= -(D_y E_z^{n+1})_{i,j+\frac{1}{2},k+\frac{1}{2}} + (D_z E_y^{n+1})_{i,j+\frac{1}{2},k+\frac{1}{2}} \\
 (D_t B_{y_{i+\frac{1}{2},j,k+\frac{1}{2}}})^{n+1} &= (D_x E_z^{n+1})_{i+\frac{1}{2},j,k+\frac{1}{2}} - (D_z E_x^{n+1})_{i+\frac{1}{2},j,k+\frac{1}{2}} \\
 (D_t B_{z_{i+\frac{1}{2},j+\frac{1}{2},k}})^{n+1} &= -(D_x E_y^{n+1})_{i+\frac{1}{2},j+\frac{1}{2},k} + (D_y E_x^{n+1})_{i+\frac{1}{2},j+\frac{1}{2},k}
 \end{aligned} \tag{IV.15}$$

where the derivative operators D_t, D_x, D_y, D_z on any given field quantity $F_{i,j,k}^n$ are defined as follows:

$$\begin{aligned}
 (D_t F)^n &= \frac{F^{n+\frac{1}{2}} - F^{n-\frac{1}{2}}}{\Delta t} \\
 (D_x F)_{i,j,k} &= \frac{F_{i+\frac{1}{2},j,k} - F_{i-\frac{1}{2},j,k}}{\Delta x} \\
 (D_y F)_{i,j,k} &= \frac{F_{i,j+\frac{1}{2},k} - F_{i,j-\frac{1}{2},k}}{\Delta y} \\
 (D_z F)_{i,j,k} &= \frac{F_{i,j,k+\frac{1}{2}} - F_{i,j,k-\frac{1}{2}}}{\Delta z}
 \end{aligned} \tag{IV.16}$$

By staggering the electric and magnetic fields both spatially and temporally as defined in the Yee lattice, the numerical error induced by the spatial and temporal derivatives over a single time step is **of order two** in both time and space:

$$\epsilon = \mathcal{O}(\Delta x)^2 + \mathcal{O}(\Delta y)^2 + \mathcal{O}(\Delta z)^2 + \mathcal{O}(\Delta t)^2 \tag{IV.17}$$

Using the FDTD scheme along with a charge conserving algorithm for the current deposition (such as the Esirkepov algorithm) guarantees that the Maxwell-Gauss equation is verified at all time steps, provided the initial electric field verifies the Maxwell-Gauss equation. Similarly, and under the same condition, the Maxwell-Thomson $\vec{\nabla} \cdot \vec{B} = 0$ is verified at all time steps:

$$\begin{aligned}
 (D_x E_x^n) + (D_y E_y^n) + (D_z E_z^n) &= \frac{\rho^n}{\epsilon_0} \\
 (D_x B_x^{n+\frac{1}{2}}) + (D_y B_y^{n+\frac{1}{2}}) + (D_z B_z^{n+\frac{1}{2}}) &= 0
 \end{aligned} \tag{IV.18}$$

Parallelization of the PIC algorithm with the FDTD scheme

The FDTD method is a relatively simple and straightforward numerical approach for solving Maxwell's equations. Moreover, the FDTD scheme is well-suited to be parallelized at arbitrary scales on distributed memory architectures. Indeed, parallelizing FDTD scheme can be achieved by employing a standard domain decomposition techniques (DD) on the simulation domain. This parallelization

strategy consists in dividing the simulation domain into small subdomains. Each subdomain is associated to a different processing unit. Neighboring subdomains only overlap over a small region called guard-region.

The computations underlying the FDTD scheme at each mesh grid point only involve data from close neighboring mesh grid cells/nodes. Therefore, the computations within different subdomains are performed independently by each processing unit. At each time step, data from the guard region intervenes to update electromagnetic quantities near the subdomains borders as shown in figure IV.3.

This parallelization strategy can be used in PIC codes and extended to parallelize the whole PIC algorithm (cf figure IV.3). Therefore, each processing unit also deals with the computations pertaining to the macro-particles that lay within its subdomain at each time step (current deposition, particles motion ...). Macro-particles are exchanged between neighboring subdomains when they cross the subdomains boundaries. Note that when using extended shape functions (with large supports), the macro-particles laying near subdomain boundaries may contribute to the current/charge of neighboring subdomains. For this reason, a different current/charge densities data exchanges is performed in order to take into account the contribution of macro-particles to external subdomains meshes (cf panel (c)). Moreover, the number of guardcells (the number of cells inside the guard region) should take into account the shape functions used in the PIC loop.

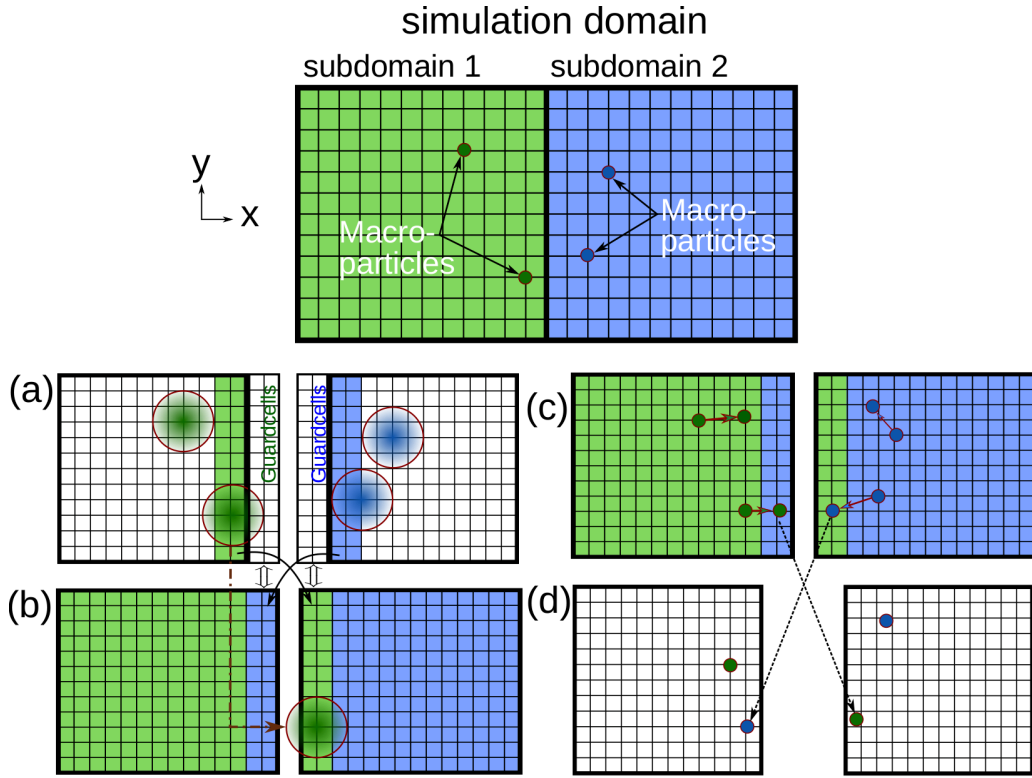


Figure IV.3: Parallelization of the PIC algorithm. The simulation domain is split into subdomains. Guardcells are appended to each simulation domain. Macro-particles deposit current and charge on the subdomain+guard region before Maxwell's equations solve (panel a). Guardcells are filled with data from neighboring subdomains (panel b). Particles that are initially close to a subdomain boundary may deposit current on a neighboring subdomain. Maxwell's equations are then solved on each subdomain. Macro-particles are pushed (panel c). If a macro-particle enters the guard region, it is sent to the corresponding subdomain. Each MPI task handles the computations associated with a unique subdomain.

This parallelization strategy is used in many PIC softwares. In the WARP+PICSAR code, this parallelization is handled by the Message-Passing-Interface library (MPI). Hence, each subdomain is associated to one MPI task.

Besides, PICSAR enables an additional shared-memory parallelization level within each MPI task: each subdomain is further subdivided into multiple smaller subdomains denoted by 'tiles'. The computations pertaining to macro-particles (field interpolation, particles push, current/charge deposition) are performed on the tile level before being merged on the MPI task mesh grid level to solve Maxwell's equations. The workload associated to tiles is distributed among different threads belonging to each MPI task, using the shared memory parallelization API OpenMP. The size of data laying within each tile is preferably set so that it fits into the fastest cache memory layer to enable efficient cache reuse and fast computations. Finally, computations within each tile are boosted by using vectorized implementations of different PIC steps via Single Instructions Multiple Data instruction set [63].

Numerical dispersion and stability condition for FDTD

The FDTD scheme has many advantages, among which its efficient parallelization over distributed memory machines. Nevertheless, it has as well various numerical limitations that hinder its accuracy.

These limitations are related to the unphysical electromagnetic field behavior when propagated via the FDTD scheme. In the context of ultra-high intensity physics PIC simulations, these drawbacks can sometimes strongly spoil the accuracy of the results obtained from PIC simulation when using the FDTD method [23].

To highlight those limitations, we derive the numerical dispersion relation associated with the FDTD scheme.

From equation IV.15, one can deduce the discretized numerical propagation equation that applies to all the electromagnetic fields components in vacuum. Considering the E_x field component for example, this leads to:

$$\Delta_x^d E_x^n + \Delta_y^d E_x^n + \Delta_z^d E_x^n = \frac{1}{c^2} \Delta_t^d E_x^n \quad (\text{IV.19})$$

where, for a given field F , the discrete second order derivative operators Δ_x^d , Δ_y^d , Δ_z^d and Δ_t^d read:

$$\begin{aligned} \Delta_x^d F_{i,j,k} &= \frac{F_{i+1,j,k} - 2F_{i,j,k} + F_{i-1,j,k}}{\Delta x^2} \\ \Delta_y^d F_{i,j,k} &= \frac{F_{i,j+1,k} - 2F_{i,j,k} + F_{i,j-1,k}}{\Delta y^2} \\ \Delta_z^d F_{i,j,k} &= \frac{F_{i,j,k+1} - 2F_{i,j,k} + F_{i,j,k-1}}{\Delta z^2} \\ \Delta_n^d F^n &= \frac{F^{n+1} - 2F^n + F^{n-1}}{\Delta t^2} \end{aligned} \quad (\text{IV.20})$$

The discrete second order derivatives operators correspond to a three point convolution in the real space. Their dual operators in Fourier space write:

$$\widehat{\Delta_\alpha^d F_\alpha} = \frac{2}{\Delta \alpha^2} (\cos(k_\alpha \Delta \alpha) - 1) \widehat{F}_\alpha = -\frac{4}{\Delta \alpha^2} \sin^2\left(\frac{k_\alpha \Delta \alpha}{2}\right) \widehat{F}_\alpha \quad (\text{IV.21})$$

where α stands for any dimension, both spatially or temporally. By Fourier transforming equation IV.19 over time and space, we retrieve the relationship between the temporal frequency ω and spatial frequencies (k_x, k_y, k_z) :

$$\frac{1}{c^2 \Delta t^2} \sin^2\left(\frac{\omega \Delta t}{2}\right) = \frac{1}{\Delta x^2} \sin^2\left(\frac{k_x \Delta x}{2}\right) + \frac{1}{\Delta y^2} \sin^2\left(\frac{k_y \Delta y}{2}\right) + \frac{1}{\Delta z^2} \sin^2\left(\frac{k_z \Delta z}{2}\right) \quad (\text{IV.22})$$

Instead of the physical dispersion relation in vacuum $\frac{\omega^2}{c^2} = k_x^2 + k_y^2 + k_z^2$, the FDTD scheme verifies a dispersion relation, given by equation IV.22. Note that when the spatio-temporal sampling goes to zero ($\omega \Delta t, k_x \Delta x, k_y \Delta y, k_z \Delta z$) $\ll 1$, the numerical dispersion relation converges to the physical one.

Another consequence of the numerical dispersion relation is the existence of a stability constraint on the temporal time step Δt , imposed by the spatial mesh sampling. For instance, when the Nyquist wave vector $\vec{K}_N = \pm \frac{\pi}{\Delta x} \vec{e}_x \pm \frac{\pi}{\Delta y} \vec{e}_y \pm \frac{\pi}{\Delta z} \vec{e}_z$ is injected into equation IV.22, we obtain:

$$\sin^2\left(\frac{\omega \Delta t}{2}\right) = c^2 \Delta t^2 \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2} \right) \quad (\text{IV.23})$$

which requires that:

$$c^2 \Delta t^2 \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2} \right) \leq 1 \quad (\text{IV.24})$$

Otherwise, ω would be a complex number, resulting in an exponentially growing instability. This condition is known as the Courant Friedrichs-Lewy (CFL) and restricts the maximum time step length in order to guarantee the stability of the scheme.

Figure IV.4, panel (a) depicts the electromagnetic phase velocity v_ϕ as a function of the wave vector (k_x, k_y) for a time step given by:

$$c^2 \Delta t^2 \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2} \right) = 1$$

The phase velocity strongly deviates from its physical value c over a large portion of the spectral domain. The amplitude of this deviation depends on the propagation angle θ in the grid. Note that the numerical phase velocity remains accurate along the line $k_x = k_y$, corresponding to electromagnetic waves propagating with a 45° angle with respect to the x-axis in the mesh grid, while it is strongly dispersive for angles closer to 0 or $\frac{\pi}{2}$ (cf panel(b)).

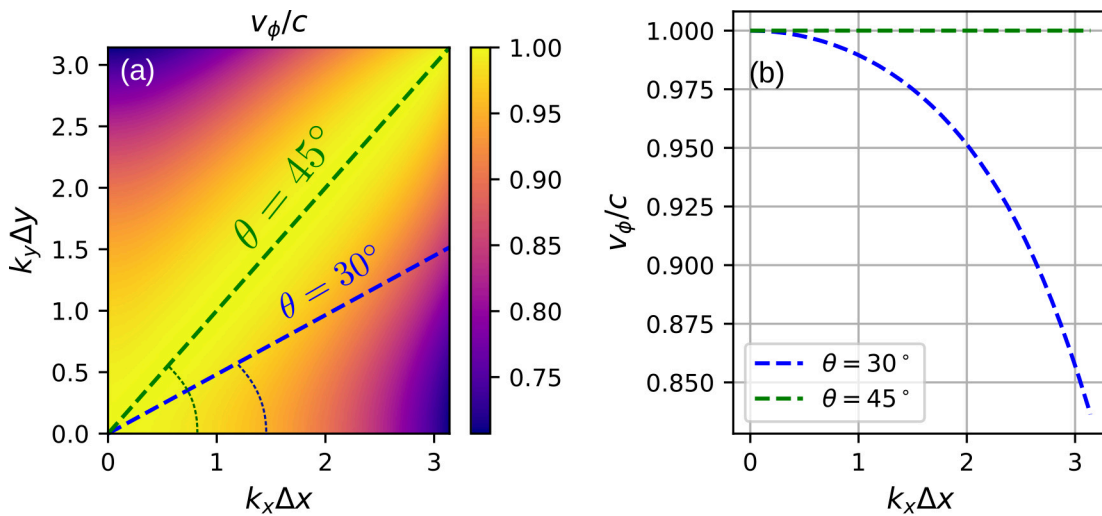


Figure IV.4: Panel (a): Numerical phase velocity in 2D as a function of (k_x, k_y) . Panel (b): Numerical phase velocity for two different angles, $\theta = 45^\circ$ and $\theta = 30^\circ$ (for $c\Delta t \sqrt{\left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}\right)} = 1$).

In many situations, the decrease in the phase velocity becomes highly detrimental to the accuracy of simulation results. It can lead, for example, to the well-known numerical Cherenkov radiation in the frame of plasma-based electron acceleration [23]. In this case, accelerated particles, traveling at a velocity larger or equal to the numerical light phase velocity emit radiations that stay in phase with the accelerated particles, leading to a resonant numerical instability.

On the other hand, in the context of high order harmonic generation on plasma mirrors, the artificial anisotropy of the FDTD scheme leads to spurious refraction of high order harmonics at the plasma vacuum interface [24].

These numerical artifacts can be avoided by employing more accurate and more sophisticated Maxwell solvers. One possible solution is to use FFT-based pseudo-spectral Maxwell solvers. This method is detailed in the next section, along with the parallelization strategies proposed in the literature.

IV.4 FFT-based Pseudo-spectral Maxwell solvers

As opposed to the FDTD scheme, which advances electromagnetic fields in the Yee lattice, pseudo-spectral methods solve Maxwell's equations in Fourier space. This type of solvers generally offers a higher level of numerical accuracy as it can be based on error-free numerical derivative computations. Nevertheless, due to the super-linear algorithmic complexity of Fast Fourier Transform (FFT) computations as well as their poor scalability efficiency on massively parallel architectures, pseudo-spectral methods have been scarcely used in Particle-In-Cell codes for a long time, where scalability and performance issues are often critical.

However, these numerical approaches have regained interest in the PIC community in recent years as a novel parallelization approach has been proposed by J-L Vay *et al* [25] that enabled an efficient weak scaling³ up to hundreds of thousands of cores [64] while overcoming numerous anomalous numerical artifacts of the FDTD solver.

WARP+PIC SAR code offers the possibility to use various formulations of pseudo-spectral FFT-based Maxwell solvers along with different parallelization strategies, some of which have been developed and implemented during this thesis.

In the following, we present the pseudo-spectral Maxwell solver formulation used for HHG simulation in this thesis. Then we discuss the parallelization strategy proposed by J-L Vay *et al* and its numerical accuracy.

IV.4.1 Pseudo-Spectral Analytical Time Domain

The Pseudo-Spectral Analytical Time Domain (PSATD) method [65] solves Maxwell's equations in the space (t, k_x, k_y, k_z) . It has the advantage of performing an analytical integration over the temporal variable, provided realistic hypotheses regarding the current and charge variations. It is, therefore, more robust than other pseudo-spectral schemes that rely on finite difference temporal integrations, such as the Pseudo-Spectral Time Domain (PSTD) method [66].

To simplify the notations, we normalize the current \hat{J} and the charge density ρ both by the vacuum permittivity ε_0 . The Maxwell's equations in the space (t, k_x, k_y, k_z) read:

$$\begin{aligned}
 \frac{\partial \hat{E}}{\partial t} &= ic^2 \vec{k} \wedge \hat{B} - \hat{J} \\
 \frac{\partial \hat{B}}{\partial t} &= -i \vec{k} \wedge \hat{E} \\
 i \vec{k} \cdot \hat{E} &= \hat{\rho} \\
 i \vec{k} \cdot \hat{B} &= 0
 \end{aligned}
 \tag{IV.25}$$

³weak scaling consists in proportionally and simultaneously increasing the amount of work and the computational resources. Algorithms that maintain a roughly constant computation time, in this case, are said to have a good weak scaling.

We define the four quantities \hat{E}^\pm and \hat{B}^\pm as following:

$$\begin{aligned}\hat{E}^\pm &= \frac{\hat{E} \pm c\tilde{k} \wedge \hat{B}}{2} \\ \hat{B}^\pm &= \frac{\hat{B} \mp \frac{\tilde{k}}{c} \wedge \hat{E}}{2}\end{aligned}\tag{IV.26}$$

with

$$\tilde{k} = \frac{\vec{k}}{|\vec{k}|}$$

By differentiating \hat{E}^\pm and \hat{B}^\pm over time we get:

$$\begin{aligned}\frac{\partial \hat{E}^\pm}{\partial t} &= \pm i|\vec{k}|\hat{E}^\pm - \frac{\hat{J}}{2} \pm ic\hat{\rho}\tilde{k} \\ \frac{\partial \hat{B}^\pm}{\partial t} &= \pm i|\vec{k}|\hat{B}^\pm \pm \frac{\tilde{k}}{c} \wedge \frac{\hat{J}}{2}\end{aligned}\tag{IV.27}$$

This set of first order ordinary differential equations [IV.27](#) can be integrated analytically using the integrating factor method for example. The electromagnetic fields \hat{E} and \hat{B} can be retrieved simply by the relationships:

$$\begin{aligned}\hat{E} &= \hat{E}^+ + \hat{E}^- \\ \hat{B} &= \hat{B}^+ + \hat{B}^-\end{aligned}\tag{IV.28}$$

This analytical integration allows to get rid of the temporal staggering between the electric and magnetic fields.

Assuming constant current and linear time dependence for the charge density over one time step, the integration of [IV.27](#) between $t = n\Delta t$ and $t = (n+1)\Delta t$ finally gives:

$$\begin{aligned}\hat{E}^{n+1} &= C\hat{E}^n + iS\tilde{k} \wedge \hat{B}^n - \frac{S}{|\vec{k}|}\hat{j}^{n+\frac{1}{2}} + i\frac{\tilde{k}}{|\vec{k}|}\left[\left(\frac{S}{c|\vec{k}|\Delta t} - 1\right)\hat{\rho}^{n+1} + \left(C - \frac{S}{c|\vec{k}|\Delta t}\right)\hat{\rho}^n\right] \\ \hat{B}^{n+1} &= C\hat{B}^n - iS\tilde{k} \wedge \hat{E}^n + i\frac{1-C}{c|\vec{k}|}\tilde{k} \wedge \hat{j}^{n+\frac{1}{2}}\end{aligned}\tag{IV.29}$$

with

$$\begin{aligned}C &= \cos(c|\vec{k}|\Delta t) \\ S &= \sin(c|\vec{k}|\Delta t)\end{aligned}$$

and where both the electric and magnetic fields are defined at integer time steps. Note also that the PSATD method applies to both the staggered and centered grids meshes. The staggered Yee grid case is handled by multiplying the spatial derivative operator in Fourier space $i\vec{k}$ by an appropriate phase factor that depends on the differentiated quantity staggering.

Thanks to the analytical integration over time and exact spatial differentiations in Fourier space, employing the set of equations [IV.29](#) to advance electromagnetic fields in vacuum rigorously preserves the vacuum dispersion relation $\omega = c|\vec{k}|$. Therefore, it is well suited for many numerical simulations scenarios where the numerical dispersion of electromagnetic waves is critical. [Figure IV.5](#) displays the expansion of a Dirac pulse using both the FDTD scheme and the PSATD algorithm. One can note

that the unphysical electromagnetic signal propagating with a phase velocity below the speed of light in the FDTD case is completely absent when using the PSATD algorithm.




Figure IV.5: Propagation of a Dirac pulse with the PSATD solver (panel (a)) and FDTD (panel(b)).

Additionally, the PSATD method imposes no CFL condition on the time step size Δt and is unconditionally stable. However, in the frame of the PIC method, the condition:

$$\Delta t \leq \frac{\min(\Delta x, \Delta y, \Delta z)}{c}$$

is imposed so that ultra-relativistic macro-particles do not travel more than one mesh cell over a single time step. This condition is necessary for the Esirkepov method.

Thus, PIC simulations employing the PSATD solver can be performed with larger time steps (and less number of iterations) compared to a similar simulation employing FDTD.

Limitations of the PSATD algorithm

It is worth to mention that using the PSATD algorithm even with the Esirkepov current/charge deposition algorithm does not conserve the Maxwell-Gauss equation. This is due to the fact that the Esirkepov method verifies the numerical continuity equations with a second order spatial derivative scheme, while Maxwell's equations are solved using an exact spatial differentiation. For this reason, divergence cleaning or current corrections methods should be used in some cases where preserving Maxwell-Gauss equation is critical. WARP+PICSAR code features various numerical techniques for stability/accuracy improvement, including divergence cleaning, current correction, and digital filtering to mitigate various instabilities. Interested readers may refer to [67, 68, 69, 25, 56].

In the simulations presented in this thesis, these methods were not used as we have observed no significant difference in the simulations results obtained with or without these correction techniques.

IV.4.2 Parallelization of the pseudo-spectral methods

For a long time, pseudo-spectral methods have not been adopted for PIC simulations. This is due to the fact that these methods are not well suited for massively parallel super-computers based on distributed memory architectures. Indeed, computing FFTs a distributed dataset between a large number of processing units requires global communications between all of these processing units (and not only between neighboring processes such as in FDTD). This type of operation exhibits poor scalability passed few thousands of cores, even on modern architectures.

In 2013, J.L-Vay *et al* demonstrated that using a PSATD solver along with a standard domain decomposition (such as the one used to parallelize the PIC loop when employing the FDTD scheme), where FFT computations are performed serially over each subdomain (instead of global parallel FFTs), efficiently suppresses numerical instabilities induced by the use of the FDTD scheme in typical laser-plasma acceleration PIC simulations while still providing very accurate results (provided that relatively large guard regions are used).

This result is rather surprising and counter intuitive as one would not expect the PSATD algorithm to bring accurate results, while getting around its main computational bottleneck in such a simple way. Hence, this robustness can be justified via physical/analytical considerations regarding Maxwell's equations, and Fourier transformation.

For instance, Maxwell's equations are **linear** partial differential equations that impose a **finite propagation speed** on electromagnetic waves. Therefore, as a result of the causality principle, the standard domain decomposition method used to parallelize the PIC computations is theoretically rigorously equivalent to a direct solving over the whole subdomain, provided that the guard regions widths are larger or equal to the speed of light multiplied by the integration time. Note that this is only true considering the continuous (and not the discretized) form of Maxwell's equations. Nevertheless, PIC simulations deal with discretized electromagnetic fields, and discrete Fourier transforms (instead of continuous Fourier transforms). The discretization process renders the propagation of electromagnetic quantities via PSATD method non causal as the computation of the spatial derivative in the Fourier space corresponds to a non local operation in the real space, involving data from arbitrarily distant grid points. Therefore, the computation of the spatial derivatives is altered by the domain decomposition. This alteration is more important near the subdomains boundaries. Therefore one would expect spurious Gibbs-like oscillations arising close to the subdomains edges. However, for sufficiently large guard regions, this numerical artifact decreases as the computation of the spatial derivatives in Fourier space becomes more accurate even near the subdomains boundaries.

IV.4.3 Truncation error and ultra-high order solvers

An extensive numerical investigation [70] rigorously quantified this noise, called truncation errors. It showed that this numerical artifact remains localized near the subdomain boundaries, and that it does not build up into large scale instabilities affecting the entire simulation domain. Moreover, the

noise amplitude decays rapidly as the number of guardcells is increased. This paper also made another important suggestion, which consists in employing ultra-high order derivative stencils instead of the exact differentiation in Fourier space to compute spatial derivatives. The underlying idea is to make the spatial derivative operators more local on the grid so that the truncation error is less important. In practice, this consists in replacing the $i\vec{k} \times \cdot$ derivative operator involved in the PSATD algorithm by the Fourier transform of a high order finite-difference derivative operator in the real space. For instance, as shown by Fornberg *et al* [71], the FDTD scheme can be generalized to any arbitrary order in space by introducing additional terms associated to more distant mesh grid points in the calculation of the spatial differentiation. The generalized derivative operator of order ($p \in 2\mathbb{N}$) reads:

$$[D_x^p(F)](l) = \frac{1}{\Delta x} \sum_{i=1}^{\frac{p}{2}} C_{\frac{p}{2}}^i \left(F_{l+(i-\frac{1}{2})} - F_{l-(i-\frac{1}{2})} \right) \quad (\text{IV.30})$$

where $C_{\frac{p}{2}}^i$ are called Fornberg coefficients and are defined in [72]. For a staggered grid, these coefficients are given by:

$$C_{\frac{p}{2}}^j = \frac{(-1)^{j+1} 16^{1-\frac{p}{2}} ((p-1)!)^2}{((2j-1)!)^2 (\frac{p}{2} + j - 1)! (\frac{p}{2} - j)! (\frac{p}{2} - 1)!^2} \quad (\text{IV.31})$$

One can show via Taylor expansions, that for a sufficiently smooth function f we have:

$$\frac{df}{dx}(l\Delta x) = [D_x^p(F)](l) + \epsilon_p \quad (\text{IV.32})$$

with:

$$\begin{aligned} F_l &= f\left(\left(l + \frac{1}{2}\right)\Delta x\right) \\ \epsilon_p &= \mathcal{O}(\Delta x)^p \end{aligned} \quad (\text{IV.33})$$

Additionally, one can define the p-order wave vector k_x^p as the dual operator in Fourier space of $\frac{D_x^p}{i}$:

$$ik_x^p = \widehat{D}_x^p = i \frac{2}{\Delta x} \sum_{j=1}^{\frac{p}{2}} C_{\frac{p}{2}}^j \sin\left(\left(j - \frac{1}{2}\right)k_x \Delta x\right) \quad (\text{IV.34})$$

This high order spatial differentiation method involves $\frac{p}{2}$ neighboring data points from each side. Therefore, it is 'more local' in the real domain than the exact operation $i\vec{k} \times$. As a consequence, the use of a finite yet high spatial derivative stencil significantly mitigates the truncation error, which completely vanishes when the number of guardcells n_g is equal to $n_g = \frac{p}{2}$. However, it is already reasonably low even when $n_g \ll \frac{p}{2}$ (see [70]).

Employing the p-order stencil derivative alongside with the PSATD algorithm guarantees a very low discretization error (cf equation IV.33) and ensures a physical numerical dispersion of electromagnetic waves. For instance, the numerical dispersion relation of the p-order PSATD algorithm reads:

$$\frac{\omega}{c} = \sqrt{(k_x^p)^2 + (k_y^p)^2 + (k_z^p)^2} \quad (\text{IV.35})$$

Panel (a) of figure IV.6 sketches the ratio $\frac{v_\phi}{c}$ for an order 100 PSATD solver. The numerical dispersion of the algorithm is very accurate over a large extent of the (k_x, k_y) space. It deviates from

the exact vacuum dispersion relation only near the Nyquist frequency $k_{\text{Nyquist}} = \frac{\pi}{\Delta x}$.

It is also important to mention that the truncation error has only a small effect on the numerical dispersion relation accuracy. In panel (b), we sketched the phase velocity $\frac{v_\phi}{c}$ obtained from an order 100 PSATD solver when the Fornberg coefficients $C_{\frac{p}{2}}^j = \frac{100}{2}$ are forced to zero beyond the 10th coefficient ($C_{\frac{p}{2}}^j = \frac{100}{2} \quad \forall j > 10$). This situation corresponds to a modified order-100 PSATD solver with 10 guardcells surrounding subdomains on each side, and where only the contributions of the 10 nearest grid points (from both sides) are taken into account in calculating the high order spatial derivatives. In practice, we do not set the Fornberg coefficients beyond n_g to zero in simulations. However, this setup is the worst case scenario that would give rise to a stencil truncation error from all grid points (even those far from the subdomains boundaries). Hopefully, even in this case, panels (b) and (c) of figure IV.6 show that the solver dispersion relation remains very accurate.




Figure IV.6: Panel (a), numerical phase velocity of the PSATD solver with an order 100 stencil. Panel (b), numerical phase velocity of the PSATD solver with an order 100 stencil but beyond the 10th coefficient. Panel (c), comparison between the phase velocities of the truncated and untruncated PSATD solvers at 45° degree.

To illustrate the accuracy of this method, we performed a series of Dirac pulse propagation simulations, using the 100 order PSATD algorithm with different numbers of guardcells (cf panel (a) of figure IV.7). We then compare the result of each test with the reference case employing 50 guardcells (where there is no error due to the stencil truncation). In all the cases, the simulation domain is split into 4×416 subdomains. As shown by panel (b) of figure IV.7, the truncation error amplitude decays exponentially with respect to the number of guardcells. And the error amplitude is already extremely

low for $n_g = 10$ (error $< 0.1\%$).

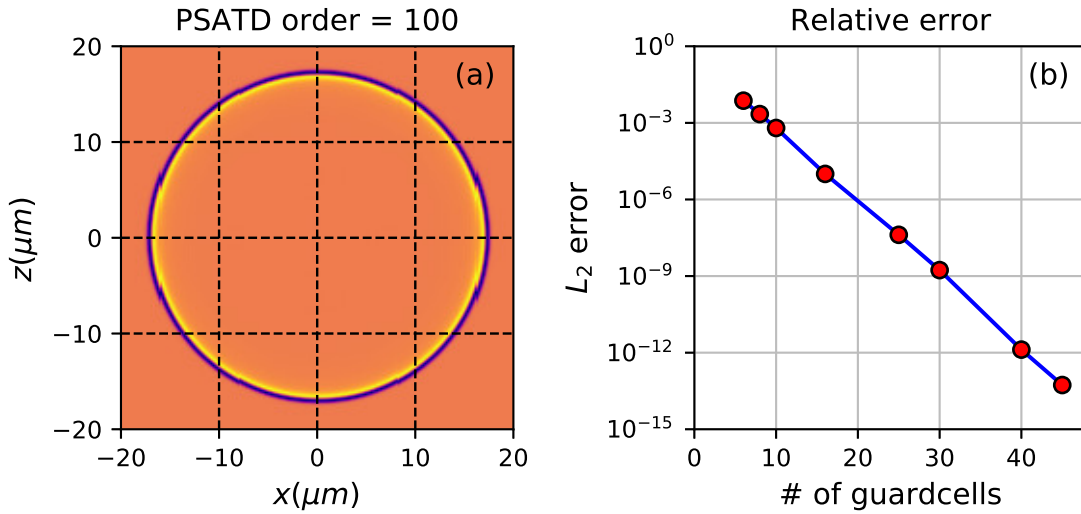


Figure IV.7: Panel (a) Propagation of a Dirac pulse with the PSATD solver and order $p=100$ stencil using 4×4 subdomains and 10 guardcells in each direction. Panel (b): Relative L_2 truncation errors as a function of the number of guardcells.

With WARP+PICSAR, the PSATD Maxwell solver can be used along with a standard domain decomposition for highly accurate and scalable PIC simulations. This implementation demonstrated an excellent weak scaling up to hundreds of thousands of cores [63]. Nevertheless, by employing relatively large guard regions, this standard parallelization strategy (called local PSATD later on) induces important data redundancy due to multiple copies of fields values stored in the guardcells of different processing units, which is detrimental to the strong scaling⁴ efficiency at very large scales as the workload pertaining to each processing unit is no longer inversely proportional to the total processing units number. In chapter V we address this limitation by proposing a novel parallelization strategy that significantly increases the scaling efficiency of the local PSATD algorithm at very large scales while substantially reducing its memory footprint.

IV.5 Absorbing boundary conditions with PSATD

To emulate open boundary conditions for electromagnetic fields, we use the Berenger Perfectly Matched Layers (PML) technique [73]. This method is based on introducing absorbing layers around the simulation domain, where electromagnetic field components are split into two parts, corresponding to different transverse electric and transverse magnetic modes along each direction (cf figure IV.8). Inside each PML region, a modified version of Maxwell's equations is solved, taking into account

⁴strong scaling consists in increasing the computational resources while keeping the amount of work constant. Algorithms are said to have a good strong scaling if the computation time scales as the inverse of the computational resources

electromagnetic absorbing coefficients :

$$\begin{aligned}
 \frac{\partial \hat{\mathcal{E}}}{\partial t} &= c^2 i \mathcal{K} \times \hat{B} - \hat{\sigma} \otimes \hat{\mathcal{E}} - \hat{J} \\
 \frac{\partial \hat{B}}{\partial t} &= -i \mathcal{K} \times \hat{E} - \hat{\sigma} \otimes \hat{B}
 \end{aligned}$$

$$\hat{\mathcal{E}} = \begin{bmatrix} \hat{E}_{xy} \\ \hat{E}_{xz} \\ \hat{E}_{yx} \\ \hat{E}_{yz} \\ \hat{E}_{zx} \\ \hat{E}_{zy} \end{bmatrix}, \hat{B} = \begin{bmatrix} \hat{B}_{xy} \\ \hat{B}_{xz} \\ \hat{B}_{yx} \\ \hat{B}_{yz} \\ \hat{B}_{zx} \\ \hat{B}_{zy} \end{bmatrix}, \hat{J} = \begin{bmatrix} \hat{J}_x \\ 0 \\ \hat{J}_y \\ 0 \\ \hat{J}_z \\ 0 \end{bmatrix}, \hat{E} = \begin{bmatrix} \hat{E}_{xy} + \hat{E}_{xz} \\ \hat{E}_{yx} + \hat{E}_{yz} \\ \hat{E}_{zx} + \hat{E}_{zy} \end{bmatrix}, \hat{B} = \begin{bmatrix} \hat{B}_{xy} + \hat{B}_{xz} \\ \hat{B}_{yx} + \hat{B}_{yz} \\ \hat{B}_{zx} + \hat{B}_{zy} \end{bmatrix} \quad (\text{IV.36})$$

$$\mathcal{K} = \begin{bmatrix} 0 & 0 & k_y \\ 0 & -k_z & 0 \\ 0 & 0 & -k_x \\ k_z & 0 & 0 \\ 0 & k_x & 0 \\ -k_y & 0 & 0 \end{bmatrix}, \quad \sigma = \begin{bmatrix} \sigma_y & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_z & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_x & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_z & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_x & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_y \end{bmatrix}$$

where $\sigma_x, \sigma_y, \sigma_z$ are electric conductivities, \otimes is the tensor product operator, and \times is matrix-vector product. Note that when these coefficients are equal to zero, equation IV.36 yields to standard Maxwell's equations.

Besides, in order to avoid numerical instabilities inside the PML region, a null current and charge are required. To enforce this condition, a macro-particles free buffer is placed between the simulation domain and the PML zone. Macro-particles reaching this buffer zone are either reflected or deleted (depending on macro-particles boundary conditions). The thickness of this buffer zone needs to be equal or larger than the deposition shape functions support to avoid any spurious currents or charge densities inside the PML region.

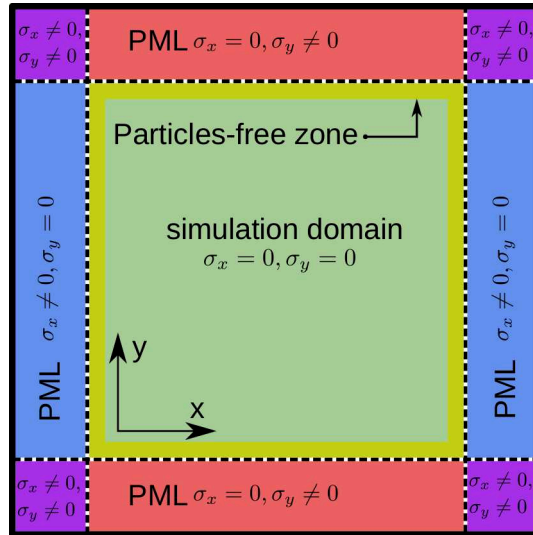


Figure IV.8: Disposition of the PML medium within the simulation box. The macro-particles free zone reflects or deletes incoming macro-particles.

Equation IV.36 holds true in the simulation domain (where the absorbing coefficients are null) and the PML zone. It can be numerically solved by employing an operator splitting methods [74] to

decouple the propagation operator $\propto i\mathcal{K} \times . - \hat{\mathcal{J}}$ from the damping operator $\hat{\sigma} \otimes$ [75]. The propagation operation consists in solving the next ordinary differential equations:

$$\begin{aligned}\frac{\partial \hat{\mathcal{E}}}{\partial t} &= c^2 i\mathcal{K} \times \hat{B} - \hat{\mathcal{J}} \\ \frac{\partial \hat{\mathcal{B}}}{\partial t} &= -i\mathcal{K} \times \hat{E}\end{aligned}\tag{IV.37}$$

This equation can be solved over the whole simulation domain (including PML regions) by employing the PSATD algorithm similarly to the formulation given in IV.29.

The damping operation is carried out in the real space via exponential temporal integration:

$$\begin{aligned}\frac{\partial \mathcal{E}}{\partial t} &= -\sigma \mathcal{E} \Rightarrow \mathcal{E}^{n+1} = \exp[-\Delta t \sigma] \mathcal{E}^n \\ \frac{\partial \mathcal{B}}{\partial t} &= -\sigma \mathcal{B} \Rightarrow \mathcal{B}^{n+1} = \exp[-\Delta t \sigma] \mathcal{B}^n\end{aligned}\tag{IV.38}$$

The exponential integration method has the advantage to be unconditionally stable for any spatial time step (as $\exp[-\Delta t \sigma] \leq 1 \quad \forall \Delta t > 0$ since $\sigma \geq 0$). Therefore, the PSATD+PML is still unconditionally stable (no CFL condition).

In conclusion, the PML method can be used alongside with PSATD for an accurate solving of Maxwell's equations under open boundary conditions constraint.

Part B

Development of numerical tools for accurate ultra high intensity physics simulations

Chapter V

A generalized massively parallel scheme for FFT-based pseudo-spectral Maxwell solvers

V.1	Introduction	71
V.2	Brief description of HPC systems	72
V.3	Scalability limits of existing parallelization techniques for pseudo-spectral solvers	73
V.3.1	Scalability limits of the global method	73
V.3.2	Scalability limits of the local method	77
V.4	A generalized massively parallel FFT-based Maxwell solver	79
V.4.1	Principle of the technique	79
V.4.2	Advantages of the new parallelization strategy	80
V.4.3	Coupling of the hybrid solver with the full PIC algorithm: a dual grid decomposition for efficient load balancing	86
V.4.4	Benchmarks of the hybrid solver on different architectures	89
V.5	Conclusion and prospects	93

V.1 Introduction

FFT-based pseudo-spectral Maxwell solvers are robust numerical tools for modelling the evolutions of electromagnetic fields in time and space with an excellent accuracy. These methods are of great interest for 3D PIC UHI simulations as they induce no numerical dispersion that would hinder the core physical phenomena at play.

Up until recently, these solvers have only been scarcely used to conduct massively parallel simulations due to the poor scalability of this method, owing to global communications associated with global FFTs computations on the entire simulation domain. Indeed, computing massively parallel FFTs on distributed-memory machines is known to be an extremely challenging task to scale, as it requires heavy and computationally expensive collective communications involving all processing units. For this reason, up to now, the scaling of global FFTs has been limited to a few tens of thousands of cores at best [76], which is not sufficient to take advantage of massively parallel super-computers full capabilities, which can reach millions of cores.

To address this scalability barrier and enable massively parallel simulations using FFT-based pseudo-spectral Maxwell solver, an important breakthrough has recently been made by J.L.Vay *et al* [25] (c.f section IV.4.1). The authors show that pseudo-spectral Maxwell solvers can be used along with a standard domain decomposition, provided that large guard regions are appended at subdomains boundaries. This way, Maxwell's equations are solved independently on each subdomain using single-node FFT computations (instead of distributed-memory FFTs), and guardcells are exchanged between adjacent subdomains at each time step. This technique, however, introduces a small numerical truncation error at the level of subdomains boundaries. In practice, the truncation errors amplitude can be efficiently mitigated either by employing a finite, yet arbitrarily high derivative stencil order, or by increasing the number of guardcells [70].

By taking advantage of the standard domain decomposition method, this parallelization strategy exhibits a very good scaling up to hundreds of thousands of cores [64] for a moderate number of guardcells ($n_g = 8$).

Nevertheless, when the number of guardcells is increased (to keep the truncation error small for ultra-high orders $p > 100$), this parallelization approach may induce an important memory footprint due to the significant data redundancy of grid arrays inside guard regions. This memory overhead also grows considerably as the number of processing units. Consequently, the strong scaling of the pseudo-spectral solver is severely impacted.

These limitations, both in terms of scaling efficiency and memory use call for a new arbitrary scalable, robust parallelization strategy that preserves good scaling performances even at very large scales and for large guardcells numbers.

In this chapter, we expose a new parallelization strategy for the pseudo-spectral Maxwell solver that outperforms previous parallelization techniques both in terms of memory footprint and scalability.

This approach introduces a coarse Cartesian domain decomposition, where each subdomain is assigned to multiple MPI tasks. Maxwell's equations are solved within this new domain decomposition, and FFTs are computed using distributed memory FFT kernels.

We first briefly describe the super-computing systems used to perform the scalability benchmarks of this chapter. Then, we highlight in detail the limitations of the parallelization approaches in use so far for the pseudo-spectral Maxwell solvers. After exposing the core ideas behind our novel approach, we show that it brings a substantial speed-up and memory gain against both the local and global parallelization (based on global FFTs) strategies. A scaling toy model for our novel scheme is presented. It allows fetching for the optimal decomposition setup that results in the best performances. We then explain how to efficiently couple this new parallelization method within a standard PIC code.

Finally, we expose the scalability results of our new solver, obtained from large scale benchmark tests performed on massively parallel ALCF computational resources, and compare them to the local pseudo-spectral solver performances.

V.2 Brief description of HPC systems

The scalability benchmarks presented in this chapter were performed on two different HPC systems: Theta-Cray XC40 ¹ and Mira-BlueGene/Q ².

- The Theta cluster is an 11 Petaflops machine based on Intel Xeon Phi processors (codenamed KNL). Theta is equipped with 4392 compute nodes, containing 64 cores each. Each core runs at a clock speed of 1.3 GHz. In total, the system has 281088 cores and a memory of 912 TB (of which 70 TB is high-bandwidth MCDRAM). The compute nodes are connected with an Aries interconnect with a Dragonfly configuration designed by Cray.
- MIRA is a 10 Petaflops IBM Blue Gene/Q system. It consists of 49152 compute nodes distributed between 48 racks. Each compute node contains 16 Power BQC 16C cores, running at a clock speed of 1.6 GHz each. In aggregate, the total RAM on MIRA is 786432 GB (16 GB for each compute node). This system is equipped with a 5D torus interconnect network.

¹<https://www.alcf.anl.gov/Theta>

²<https://www.alcf.anl.gov/Mira>

V.3 Scalability limits of existing parallelization techniques for pseudo-spectral solvers

V.3.1 Scalability limits of the global method

As already mentioned, the global parallelization method is generally not suited for massively parallel simulations as it is limited by the poor scaling efficiency of distributed-memory FFT at large scales. The next section presents the main parallelization strategies used by most modern high performance computing FFT software packages and explains why these implementations usually exhibit poor scalability performance and very large scales. This behavior has been assessed through a series of benchmarks performed on MIRA cluster at ALCF using the PICSAR code.

Since our new parallelization method for the pseudo-spectral Maxwell solver heavily relies on distributed-memory FFT, understanding FFT performance patterns and limitations will be essential to fully take advantage of the new hybrid solver.

Overview of distributed-memory FFTs

Handling the collective communications underlying distributed-memory FFT computation can be achieved using different approaches (refer to [77, 78] for more information). The most effective strategy already in use in many high performance FFT libraries is the so-called "the transpose transform". This algorithm applies to multidimensional FFTs and consists in successively alternating between computing serial 1D FFTs along the dimensions where the data resides entirely in a single processor's memory and remapping the data to "localize" the initially distributed axes. The transposition step is essentially a global operation that involves all-to-all communication patterns. This type of operation is usually poorly scalable at large core count due to the growing communication contention and network saturation.

A schematic illustration of this algorithm is shown in figure V.1 and can be summarized as follows:

1. Grid array is decomposed into sub-arrays along one or multiple axes. Each processing unit stores one sub-array.
2. Perform single node 1D FFT along each axis for which data resides entirely in each processor's memory. This computation is performed by each processor independently and can be achieved fairly fast.
3. Select one of the remaining axes along which 1D FFT that has not been computed yet. Remap the data array decomposition to "localize" the data along the new axis on each processor and enable step 1. This step is usually performed using `MPLAlltoall` routines.
4. Repeat step 1 and step 2 until all axes are processed.
5. Finally, rearrange the output array so that it has the same memory layout as the input array.

This step is usually optional and can be omitted to save computation time.

This scheme is implemented by many FFT libraries [79, 80, 81] but it has only demonstrated -at best- a good scaling up to a few tens of thousands of cores only. Historically, the first implementations of distributed-memory FFTs using the transpose scheme relied on 1D domain decomposition. This scheme is known as the "slab decomposition" (see figure V.1). A multidimensional array is distributed between different computing cores along a single axis.

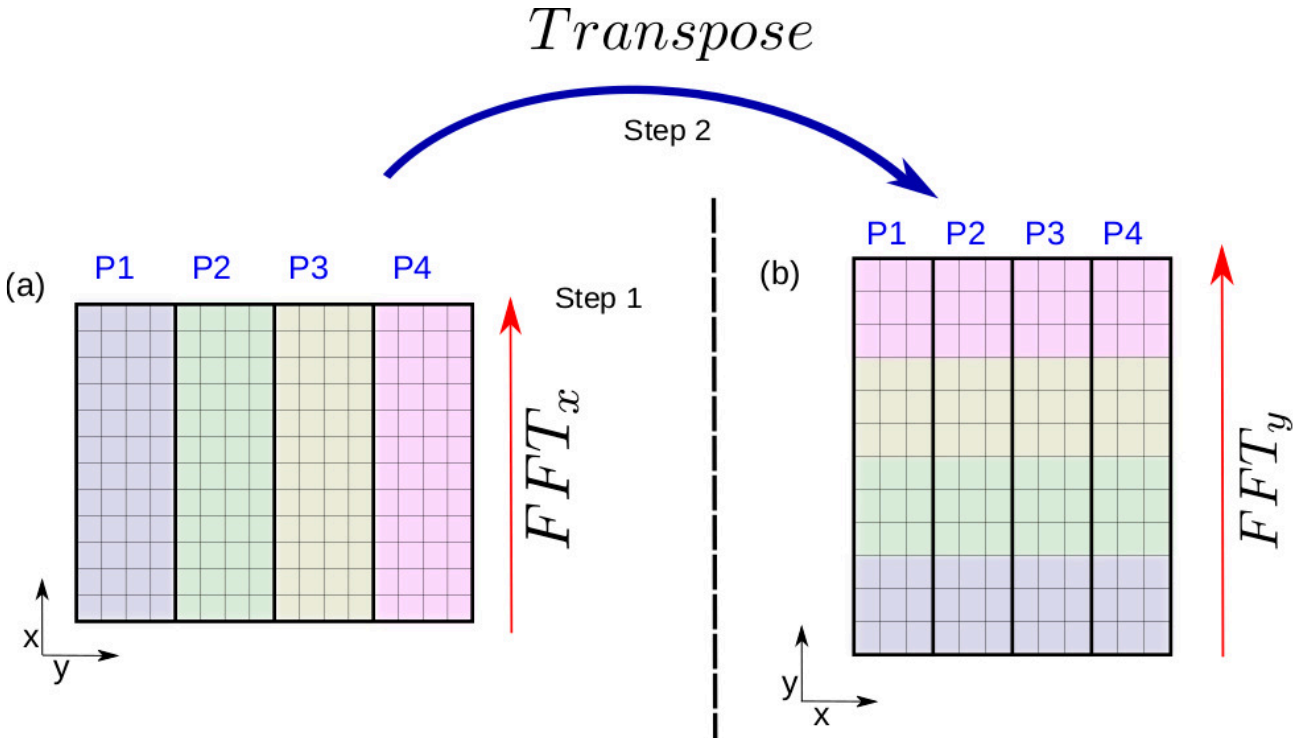


Figure V.1: Schematic representation of parallel FFT computation using the slab decomposition. Each processor stores one chunk of data corresponding to a "slab" that is orthogonal to the y -axis. This initial memory layout allows computing FFT along x -direction serially. The transposition permutes x and y axes so that each slab is now parallel to the y -axis. This memory layout allows one last 1D FFT computation along the y -axis.

In practice, the partitioned axis is usually chosen to be the axis along which data points are the farthest in memory (first axis for C arrays, last axis for Fortran arrays).

This technique allows performing distributed memory-FFT with only one transposition. Due to its simplicity, it is still in use by many distributed-FFT software packages, including the well-known open-source FFTW library [82] as well as the Intel-provided MKL-FFT package. However, since the data decomposition is only performed along one axis, this technique cannot be used whenever the number of processes exceeds the number of data points along the split axis. This limitation is very constraining especially when dealing with 3D data arrays on massively parallel machines (involving hundreds of thousands of cores).

More flexible approaches have been developed to allow the use of a larger number of processing units. Nowadays, many FFT libraries employ a 2D data decomposition to perform n -dimensional FFTs (with $n \geq 3D$). This technique is known as the "pencil decomposition" and is illustrated by

figure V.2: a 3D grid array is partitioned along two dimensions. Only FFT computations along the x-axis can be performed serially, while additional transpositions are required to compute FFT along the partitioned axes. This technique is used by the well-known P3DFFT [80]. The pencil decomposition overcomes the scalability barrier of the slab decomposition and is usually sufficient for large scale 3D problems, but it requires two data transpositions in order to deal with the distributed axes.

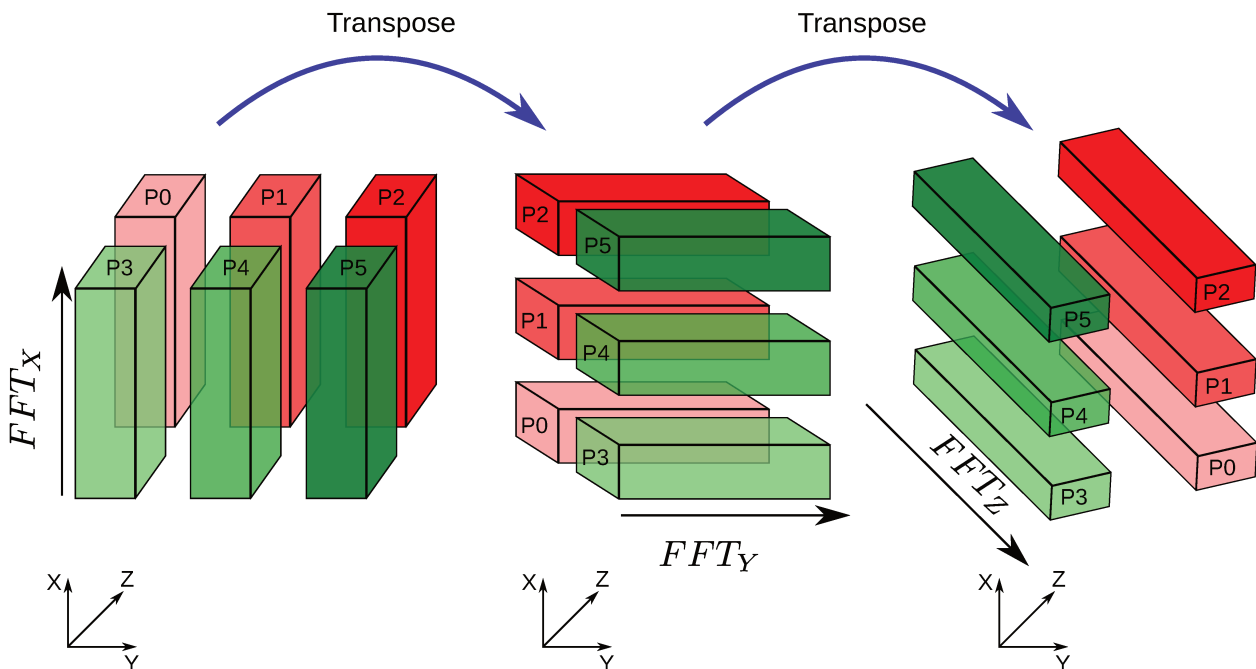


Figure V.2: 3D FFT performed with a pencil decomposition using 3 by 2 processor grid. Initially, (left sketch) each processor's stores one chunk of data corresponding to a "pencil" that is parallel to the x-direction. This memory layout allows to perform a 1D FFT along the x-axis. The first transposition permutes the localized x-axis with a distributed one (y-axis) and allows performing 1D FFT along the y-axis (central sketch). Finally, the data is transposed again to perform one last 1D FFT along z-axis (right sketch). Note that in the end, the data layout in Fourier space is not the same as in the real space, and optional reordering, involving an additional global transposition can be performed to have the same layout for real and complex data.

Note that more recent high performance computing packages support higher dimensions domain decomposition: the Plimpton package for example³, is a 2D-3D FFT library that supports all input data layouts, including 3D data decomposition for 3D FFTs. Similar techniques have been reported to be in use to perform parallel 3D FFTs in molecular dynamics simulations [83]. These codes can operate distributed-memory FFTs even when no axis is initially localized in memory and offer more flexibility regarding initial data layout. However, this approach is not always optimal to minimize computation time. For instance, handling the additional decomposed axis usually requires internal data reshuffling to remap the data into a pencil-like setup, inducing extra time and communications overheads, before performing computations following the pencil-FFT strategy.

³<https://www.sandia.gov/sjplimp/docs/fft/README.html>

Scaling of the global pseudo-spectral Maxwell solver

The global parallelization approach of the pseudo-spectral solver is usually not an efficient strategy to tackle massively parallel 3D UHI-PIC simulations on current and future super-computers that may require millions of cores ⁴.

To assess the scalability limits of the global pseudo-spectral Maxwell solver, a series of strong scaling benchmarks have been performed on MIRA BlueGene/Q cluster, using both P3DFFT and FFTW-MPI. The benchmarks parameters are summarized in the table V.1 below:

FFT Library	n_x	n_y	n_z	n_{threads}	$n_{\text{proc_min}}$	$n_{\text{proc_max}}$
FFTW-MPI	64	128	262144	1	2^{13}	2^{18}
P3DFFT	512	2048	2048	1	2^{13}	2^{18}

Table V.1: Benchmark parameters on the MIRA cluster.

As illustrated by figure V.3, these benchmarks show a scaling efficiency of only 45% between 512 and 16384 nodes when using P3DFFT and while FFTW-MPI does not scale at all, resulting in an increase of the computation time as the number of processes is increased.

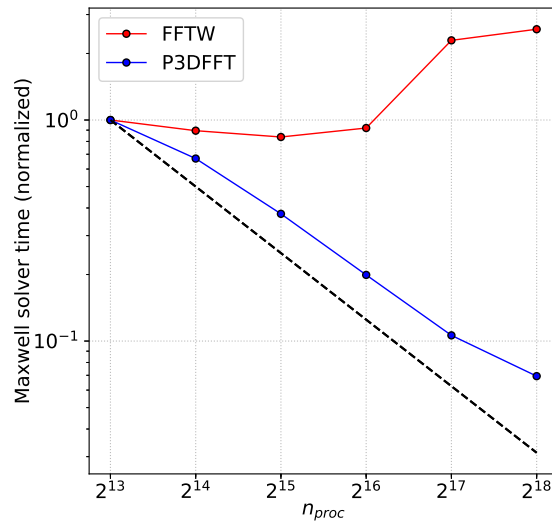


Figure V.3: Strong scaling of the global pseudo-spectral Maxwell solver using P3DFFT(blue curve) and FFTW-MPI(red curve) on the MIRA cluster. The black dashed curve represents a perfect scaling.

In conclusion, the global parallelization scheme brings an unmatched accuracy in solving Maxwell's equations as it introduces no spurious numerical noise from stencil truncation. Nevertheless, due to poor scalability performances of distributed-memory FFT algorithms on massively parallel super-computers, this parallelization scheme cannot be used for massively parallel 3D PIC-UHI simulations.

⁴www.top500.org

V.3.2 Scalability limits of the local method

By employing relatively large guard regions, the local pseudo-spectral Maxwell solver induces important data redundancy due to multiple copies of fields values stored in the guardcells of different processing units. Thus, the memory footprint of the local solver may experience a significant growth simply by increasing the number of processing units or the number of guardcells in the simulation. Owing to the fact that FFT computations are performed on the sub-arrays formed by the fields values on individual subdomains and guard regions, the scalability efficiency of the local solver is significantly mitigated in the presence of an important data redundancy. To illustrate this, let us consider a fixed simulation box of size N^d , where d is the dimension number, and for n_g guardcells, the computation time of FFTs performed within the local pseudo-spectral solver can be written as :

$$\begin{aligned} T_{fft} &= \mathcal{O}(2n_g + n)^d \log(2n_g + n) \\ &= \mathcal{O}\left(2n_g + \frac{N}{n_p}\right)^d \log\left(2n_g + \frac{N}{n_p}\right) \end{aligned} \quad (\text{V.1})$$

where n_p is the number of processes along each direction and $n = \frac{N}{n_p}$. The scaling of equation V.1 is pictured in figure V.4: FFT computations successfully scale well until a certain threshold (around 2^4 MPI tasks per dimension), where they start losing performance for larger core counts. Indeed, the total volume of data ($(2n_g + n)^d$) on which the FFT is computed does not scale with the number of processes since the number of guardcells is constant. When the volume of the guard regions becomes significant compared to the total volume of the local subdomain without guardcells (i.e. when the ratio $\frac{n_g}{n}$ increases), the FFT computation time suffers from an important scalability loss.




Figure V.4: Evolution of the FFT computation time as a function of the number of MPI tasks per axis in 1D (Magenta curve), 2D (Blue curve) and 3D (Red curve). The problem size is $N^d = 2048^d$ and $n_g = 16$. The dashed line represents a perfect scaling.

To assess this behavior in a realistic case, a massively parallel benchmark has been performed on the THETA cluster. This benchmark measures the strong scaling efficiency of the local solver at very large scales for different numbers of guardcells. The results of this benchmark are presented in figure

V.5. As expected, while the local solver succeeds in keeping a good strong scaling for a low number of guardcells (red curves of panels (a-b)), it experiences an important scalability loss at high number of guardcells as well as a substantial increase in memory volume occupied by fields quantities (panel (c)).

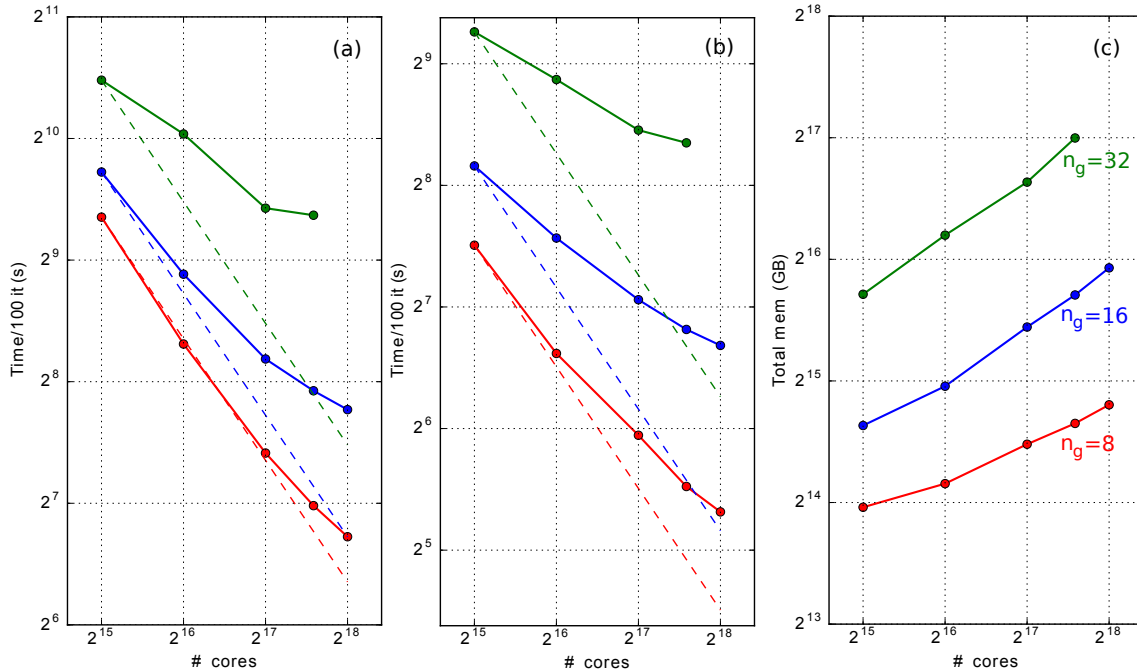


Figure V.5: Evolution of strong scaling and memory consumption (grid quantities) of the pseudo-spectral PIC algorithm for different number of guardcells n_g ($n_g = 8$: red, $1n_g = 16$: blue, $n_g = 32$: green) on Cray XC40 THETA cluster at ALCF (using 32768-262144 cores) with one OpenMP thread per MPI task. The FFTs are computed using the Intel MKL library which is more efficient on Intel architectures than FFTW. The simulation box consists in a homogeneous plasma with one particle per cell for both electrons and ions. The grid size is $240 \times 6144 \times 12288$ grid cells. Panel (a) represents the scaling of the full PIC loop. Panel (b) represents the scaling of the pseudo-spectral solver only, including MPI-exchanges for grid quantities. Panel (c) represents the total memory consumption of grid quantities (i.e. field quantities).

This shows that the local pseudo-spectral solver may not be suited for many simulations employing extremely high spatial derivative orders and requiring very large guard regions to mitigate truncation errors. In practice, multiple numerical PIC-UHI scenarios would require very large guard regions ($n_g \sim 20 - 40$) in order to allow the use of dispersion free ultra high order stencils $p \sim 250$. For example, hindering numerical Cherenkov instabilities in Plasma acceleration simulations would require a very high order stencil to ensure efficient mitigation of Numerical Cherenkov Instabilities [23] [84].

In conclusion, although local solvers are an important breakthrough towards realistic massively parallel PIC-UHI simulations, two related issues may still need to be addressed. First, the important memory overhead induced by the Cartesian domain decomposition when employing large guardcells numbers. And second, the scalability performance loss of FFT computation at very large scales. Both these limitations motivated the development of a new arbitrarily scalable parallelization strategy

allowing the use of higher number of guardcells at a higher core counts while ensuring a good scaling efficiency.

V.4 A generalized massively parallel FFT-based Maxwell solver

In this section I now present the new parallelization technique for the FFT-based Maxwell solver. This new parallelization strategy (later called "hybrid FFT-based solver") is a more flexible approach than the local solver as it allows the use of an arbitrarily high number of guardcells/processes while still ensuring an extremely good strong/weak scaling and a substantially lower memory footprint compared to the "local" parallelization. This strategy has been benchmarked on super-computers at Argonne Leadership Computing Facility. An important gain in terms of both memory ($\times 8$) and computation time ($\times 3$) has been demonstrated in these benchmarks.

V.4.1 Principle of the technique

The general principle of the new strategy is inspired from both the "local" and "global" parallelization strategies [85]: while the "global" parallelization has a very low memory footprint as there is no data redundancy due to the lack of guardcells (but scales poorly at moderate/large scales), the local strategy keeps a good scaling as the ratio $\frac{n_g}{n}$ stays small (but induces an important memory overhead at very large scales).

Figure V.6 schematically illustrates the basics of our approach. In addition to the "fine" domain decomposition in which each subdomain is assigned to one MPI task, a coarser domain decomposition is performed over the simulation domain. Each coarse subdomain (later called MPI group) gathers a cluster of MPI tasks. Guardcells are now solely appended to the MPI group boundaries. At each time step, Maxwell's equations are solved with the pseudo-spectral solver using the new domain decomposition and guardcells are exchanged between MPI groups. The FFT computations are performed using a distributed-memory FFT along all the processing units assigned to each MPI group. This technique reduces the ratio of guardcells volume to the simulation box size, which results in an important memory gain compared to the standard local solver (based on single-node FFTs) and ensures an unprecedented level of scaling efficiency at very large scales even when using large guard regions.

The choice of the group sizes (i.e. the number of MPI tasks assigned to each group) is critical to get the better performance out of this strategy: on the one hand, while very large groups would fall to the poor performance of the distributed-memory FFT, very small groups on the other hand would introduce an important memory redundancy which may as well result in a poor scaling.

The hybrid parallelization technique has been implemented in the WARP+PICSAR framework, alongside the local parallelization technique already in use. The implementation relies on FFTW-MPI or P3DFFT external libraries to perform distributed-memory FFTs computations inside each group,

with a clear performance edge for the pencil decomposition.

With fairly optimal choices of group sizes, this technique allows for a significant memory saving by reducing the data redundancy due to guardcells, and a much better strong scaling than both the local and global solvers.

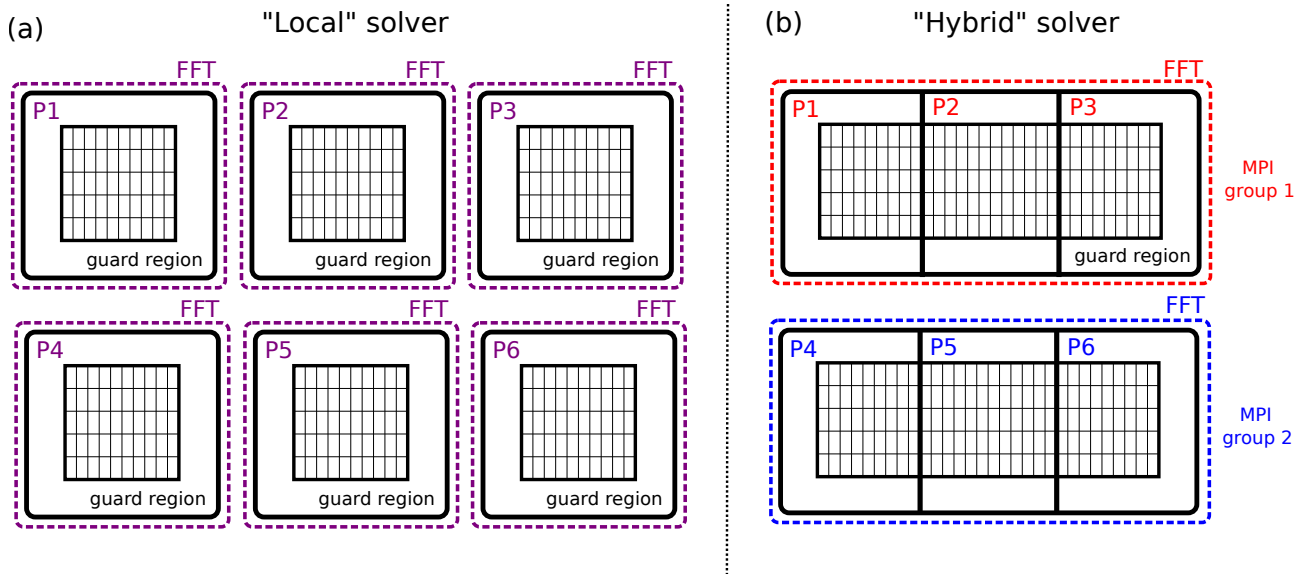


Figure V.6: Parallelization strategies for pseudo-spectral Maxwell solvers. (a) is a sketch of the 'local' approach where the simulation domain is split into multiple subdomains with guardcells appended at each subdomain boundary. Guardcells hold copies of electromagnetic fields from adjacent subdomains. Each subdomain is handled independently by each MPI process. At each time step: (i) Maxwell's equations are advanced independently on each MPI subdomain using shared-memory 'local' FFTs and (ii) guardcells are exchanged between adjacent MPI subdomains. Panel (b) shows a sketch of the new 'hybrid' approach presented in this thesis. It consists in grouping several MPI subdomains into a larger MPI group and performing distributed FFTs over the MPI tasks of this group. Guardcells are appended at the boundaries of each MPI group, leading to less memory redundancy and thus to a significant memory saving. At each time step: (i) Maxwell's equations are advanced independently on a MPI group using a distributed FFT, (ii) guardcells are exchanged between MPI groups.

V.4.2 Advantages of the new parallelization strategy

Gain in terms of Memory footprint

This hybrid strategy allows for a significant saving of data sizes by reducing the surface to volume ratio of each subdomain. To illustrate this, let us assume a cubic mesh of size $n_x \times n_y \times n_z = N^3$ split into n_p MPI subdomains in each direction x, y and z, n_g guardcells are used for each MPI subdomain, the total memory occupied by electromagnetic field arrays varies as :

$$M_{tot}^{loc} = \mathcal{O} \left((n_p^3 \times [\frac{N}{n_p} + 2n_g]^3) \right) \quad (\text{V.2})$$

The total memory occupied by field quantities strongly increases with the number of guardcells and the number of processes. The maximum number of subdomains in each axis for this problem is given by $n_p = \frac{N}{n_g}$, for which the total memory volume of electromagnetic arrays culminates at :

$$M_{tot}^{loc} = 27M_n$$

where M_n would be the total memory occupied by fields arrays without any extra memory coming from guardcells. This maximum is attained for a much lower number of processes when $n_g \gg 1$.

Let us assume that MPI domains are now grouped into MPI groups along d directions (with $d \leq 3$). Each MPI group contains n_{mpi} MPI subdomains along each gathered direction. In this case, the total memory occupied by electromagnetic field arrays reads :

$$M_{tot}^{hyb} = \mathcal{O} \frac{n_p^3}{n_{mpi}^d} \left[\frac{N}{n_p} + 2n_g \right]^{3-d} \left[n_{mpi} \frac{N}{n_p} + 2n_g \right]^d \quad (\text{V.3})$$

for $\frac{N}{n_p} = n_g$, we obtain:

$$M_{tot}^{hyb} = 3^{3-d} \left(\frac{2 + n_{mpi}}{n_{mpi}} \right)^d \times M_n$$

The memory gain of the hybrid solver compared to the local solver for the critical case where $n_p = \frac{N}{n_g}$ reads:

$$G_d^3 = 3^d \left(\frac{n_{mpi}}{2 + n_{mpi}} \right)^d \quad (\text{V.4})$$

In practice, the number d of axes along which MPI subdomains can be grouped depends on the number of axes along which the distributed-memory FFT can be parallelized, with $d = 1$ for the slab FFT distribution and $d = 2$ for the pencil FFT distribution. The larger d is, the more memory saving can be achieved. The FFTW library allows 1D slab decomposition. Thus, MPI tasks can only be grouped along z-direction (in Fortran). P3DFFT library [80], on the other hand, allows 2D pencil grouping along y and z axes in Fortran. The gain G_d^3 as a function of n_{mpi} is represented in fig V.7 for different values of d . For relatively small values of $n_{mpi} \simeq 15$, the memory gain can already approach its maximum asymptotic value ($G_d^3 = 3^d$) which can be achieved when using the global parallelization technique (that can be emulated by setting the number of groups to one).




Figure V.7: Memory gain G_d^3 of the hybrid solver compared to the local solver (corresponding to $n_{mpi} = 1$) as a function of the number of MPIs per group along each axis for different values of d , the dashed lines represent the asymptotic gain in each case.

Gains in terms of guardcells data exchanges

By appending guardcells solely to groups of MPI, the volume of exchanged data between individual MPI ranks will be reduced. The gain in terms of data exchanges can be calculated by:

$$G_{guards}^d = \frac{M_{tot}^{loc} - M_n}{M_{tot}^{hyb} - M_n}$$

For $\frac{N}{n_p} = n_g$, the previous formula becomes:

$$G_{guards}^d = \frac{3^d - 1}{\left(1 + 2/n_{mpi}\right)^d - 1}$$

The maximum gain on the total volume of guardcells exchanges for $d = 2$ (pencil decomposition with P3DFFT) and $n_{mpi} = 6$ is more than 10.

Improved scaling of FFT

The main advantage of using the hybrid parallelization strategy emerges from the improved scaling of FFT computations. In the following, we first estimate the time complexity of the distributed-memory FFT algorithm. In the light of this estimate, the advantages of hybrid solver compared to both the local and the global parallelization techniques are highlighted.

Estimating distributed-memory FFTs computation time:

Assessing the time complexity of the distributed-memory FFT is essential to understand and fully exploit the hybrid solver scaling performance at its full potential.

As explained in section V.3.1, distributed-memory FFTs is an alternation between the transposition operation and the single-node FFT computation. Based on this reasoning, one can define the total time required to perform a 3D FFT T_{fft} as the sum of the time required to perform each of the steps.

$$T_{fft} = T_c + T_{tr} \quad (\text{V.5})$$

where T_c and T_{tr} stand for the total single-node FFTs computing time and and the total transposition time operation respectively.

The time complexity of the single-node FFT computations is given by:

$$T_c = \alpha N^3 \log(N^3) \quad (\text{V.6})$$

where N is the array size along each direction, assuming a cubic array and α is a machine-dependent parameter. Assuming a pencil decomposition with $n_{proc} = n_p^2$, each processor will perform $\frac{N^3}{n_{proc}}$ 1D FFTs. Thus V.6 reads:

$$T_c = \alpha \left(\frac{N^3 \log(N^3)}{n_{proc}} \right) \quad (\text{V.7})$$

On the other hand, data transposition complexity is very network dependent. The all-to-all communications underlying the FFT computations involve each MPI task sending/receiving $n_{proc} - 1$ messages of size $\frac{N^3}{n_{proc}^2}$. Following the arguments presented in [80], the transposition time can be approximated by:

$$T_{tr} = \beta \left(\frac{N^3}{\sigma_{bi}(n_{nproc})} \right) \quad (\text{V.8})$$

where $\sigma_{bi}(n_{nproc})$ is the bisection bandwidth of the portion of the network containing n_{proc} processes and β is an architecture dependent parameter. From equations V.8 and V.7, equation V.5 can be written as :

$$T_{fft} = \alpha \left(\frac{N^3 \log(N^3)}{n_{proc}} \right) + \beta \left(\frac{N^3}{\sigma_{bi}(n_{nproc})} \right) \quad (\text{V.9})$$

To have a good estimate of the bisection bandwidth as a function of n_{proc} , one should refer to the network architecture specifications. For a 5D torus interconnect such as the one equipping the IBM BG-Q MIRA cluster at the ALCF, the bisection bandwidth scales as $n_{proc}^{4/5}$.

In the case of THETA cluster, although the dragonfly network is not an n_d torus, numerical benchmarks showed a similar asymptotic scaling for collective communication as on MIRA. As opposed to MIRA, where compute nodes are allocated contiguously, THETA compute nodes can be allocated at remote locations on the network, depending on the machine occupancy at a given time. Based on the bisection bandwidth estimates, the following formula was used for both machines:

$$T_{fft} = \alpha \left(\frac{N^3 \log(N^3)}{n_{proc}} \right) + \beta \left(\frac{N^3}{n_{proc}^{4/5}} \right) \quad (\text{V.10})$$

In the light of equation V.10, knowing the values of α and β is sufficient to properly assess the best MPI groups configuration for a given problem (assuming fixed box size, computational resources and guard region sizes).

Figure V.8 shows the behavior of the FFT computation time of the hybrid solver based on equation V.10, depending on the total number of MPI groups. We assumed that $\frac{\beta}{\alpha} = 8$ (which constitutes a realistic estimation, comparable to that measured for a Dragonfly network), a box size of $n_x = n_y = n_z = 4096$ with $n_g = 16$ and a total number of processes $n_{proc} = 128^2$. This figure shows, assuming that the scaling model of distributed memory FFTs (equation V.10) is correct, that the hybrid solver may perform better than both the local and the global solvers as it reduces the total volume of data over which FFTs are computed while keeping the transposition time overhead in an acceptable range.




Figure V.8: Theoretical FFT computation time as a function of the number of groups per direction based on equation V.10. We supposed a pencil-like parallelization with the same number of groups and processes along y and z. For the local parallelization case -corresponding to $n_{groups} = 2^7$, we have omitted the transposition term as there is no transposition involved in this case.

The choice of groups configuration strongly affects the performance of the solver: for a large number of processes per group (i.e. the global solver), the relative cost of the transposition in equation V.10 is dominant over the 1D FFT computation times, hence the distributed-memory FFT scaling is proportional to $n_{proc}^{4/5}$. On the other hand, for a low number of processes per group, the 1D FFT computation time is dominant, but FFTs are performed on larger volumes of data due to data redundancy, which also hinders the total performance of the solver. If the values of β and α are known, this toy model can be used to investigate the best MPI groups configuration for the hybrid solver by minimizing the value of T_{fft} in equation V.10.

A series of benchmarks have been performed on THETA and MIRA clusters in order to assess the accuracy of this scaling model and to investigate the best hybrid configurations on these machines. The benchmarks consisted in solving Maxwell's equations over few time steps using different numbers of guardcells and MPI groups sizes at each run. The total number of processes is kept constant on

each machine. The results of these benchmarks are sketched in figure V.9. Each panel corresponds to one cluster: the dots represent the measured FFT computation time as a function of MPI groups sizes. These sets of data are then used to estimate α and β of equation V.10 based on a least-square fit. The fitting curves are sketched with dashed lines. This set of benchmarks shows that the scaling model of distributed memory FFT accurately fits well with the measured FFT computation times: the gap between the fitting curves and the actual data does not exceed 18% on both MIRA and THETA clusters for the worst case. As expected, larger groups are more optimal when using large guard regions while smaller groups (and even local solvers in some cases) are more suited for small guard regions sizes.




Figure V.9: FFT execution time on MIRA (right panel) and THETA (left panel) as a function of the number of MPIs per group, using a pencil decomposition. The dashed lines represent the least square fit, and the dots represent the measured data. The triangles on the left panel represent the scaling of the local solver.

Benchmark on THETA:

The benchmark on THETA consisted in a 3D simulation box with the following parameters:

- $n_x \times n_y \times n_z = 512 \times 4096 \times 4096$
- $n_g = 8, 16, 32$
- $n_{proc} = 512 \times 64 = 32768$ MPI ranks (On 512 KNL nodes).
- $n_{px} = 2$ (2 processes along x)
- 16384 MPI ranks were split equally between y and z axes.

The least square fit for α and β gave:

$$\frac{\beta}{\alpha} = 4$$

This ratio reflects the relative cost of collective communications over effective computations involved in the FFT. On THETA, collective communications dominate the parallel FFT computation time, which means that this architecture favors small MPI groups.

The left panel of figure V.9 shows that the estimated scaling correctly fits the simulation time measured from the benchmarks. From this test, we can note that the local solver performs better than all hybrid parallelization for 8 guardcells (red curve), which did not come as a surprise given the value of $\frac{\beta}{\alpha}$ ratio.

Benchmark on MIRA:

The benchmark on MIRA also consisted in a 3D simulation box with the following parameters:

- $n_x \times n_y \times n_z = 2048 \times 2048 \times 2048$
- $n_g = 8, 16, 32, 64$
- $n_{proc} = 8192 \times 8 = 65536$ MPI ranks (On 8192 nodes with two OpenMP threads per MPI rank)

In this benchmark, the number of MPI ranks along the x-axis, n_{px} is also varied at each run, such that $n_{px} = n_{groupy} = n_{groupz}$, where n_{groupy} and n_{groupz} are the number of groups along y and z respectively.

The results from these tests are sketched on the right panel of figure V.9. The least square fit for α and β gave :

$$\frac{\beta}{\alpha} = 1.67$$

This value reflects that MIRA BlueGene/Q cluster is more effective in handling collective communications. Therefore, the transposition operation does not dominate that much the total FFT computation time as it does on THETA system. This can be explained by a more stable network interconnect and a processes mapping protocol that guarantees more MPI contiguity. On MIRA, optimal parallelization configurations usually correspond to large groups covering hundreds to thousands of compute nodes. Note also that it was not possible to benchmark the local parallelization scheme in this test due to memory shortage. Indeed, MIRA has much less memory per core compared to Theta (1GB per core on MIRA for 3.25 GB for Theta), and memory management is of critical importance in order to perform large scale simulations on this machine, which highlights the advantage of using our novel parallelization approach on this architecture.

V.4.3 Coupling of the hybrid solver with the full PIC algorithm: a dual grid decomposition for efficient load balancing

All the challenge in coupling the hybrid parallelization scheme with the PIC algorithm lies in the difficulty related to load balancing the particles and particles-mesh computations (particle push,

current deposition, fields interpolation) on one hand and the hybrid Maxwell Solver on the other hand. This is illustrated by panels (a) and (b) of figure V.10. Panel (a) shows the domain decomposition D_1 used to efficiently load balance particles and particles /mesh computations (assuming a homogeneous plasma distribution). Limits of MPI subdomains are highlighted using solid black lines. Guardcells required for particle-mesh operations and of width equal to the order of deposition/gathering have not been represented for more clarity. Panel (b) shows the domain decomposition D_2 that would be required to efficiently load balance the FFTs workload in the Maxwell solver step: Groups guardcells are unequally distributed between processes forming each group. Each MPI task subdomain will be shifted and/or extended due to the presence of guardcells appended at the MPI group boundaries and needed in the computation of FFTs. One can see that in that case, the limits of MPI subdomains in D_1 do not coincide with the limits of MPI subdomains in D_2 .

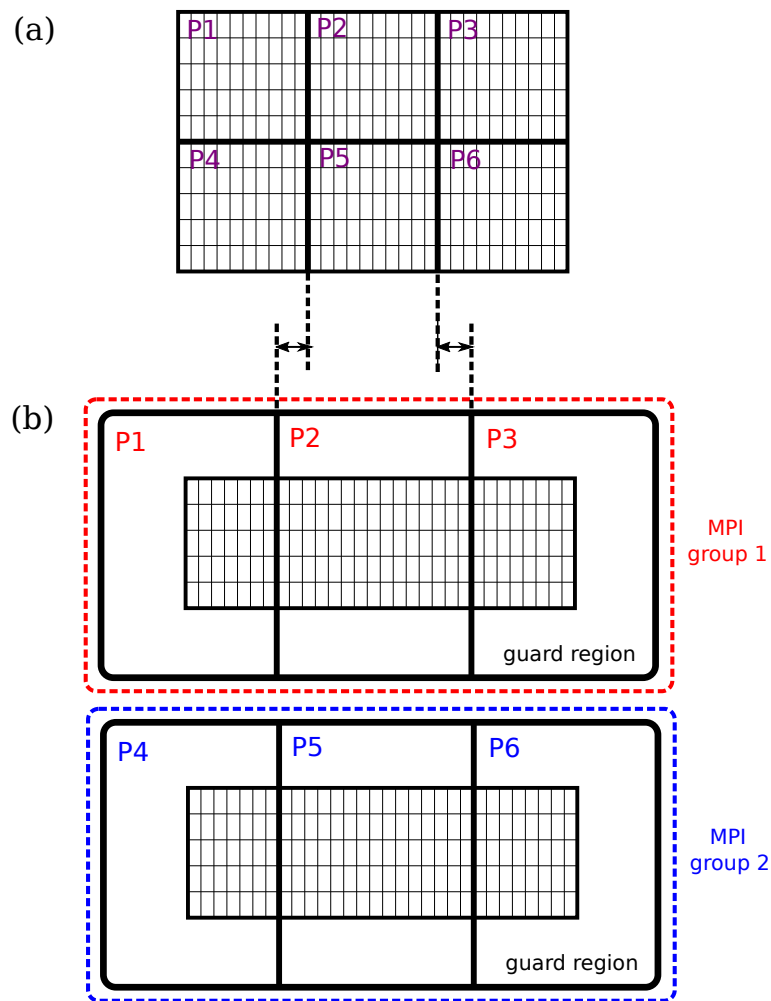


Figure V.10: Dual grid decomposition used to load balance: particles and particles/mesh computations of the PIC loop (panel a) and the hybrid solver (panel b). Note the data shift induced by the hybrid Maxwell grid. Panel (a) shows the domain decomposition D_1 used to efficiently load balance particle and particle/mesh operations (plasma is assumed to be homogeneous). Limits of MPI subdomains are highlighted using black solid lines. Guardcells required for particle-mesh operations (deposition/gathering) and of width equal to the order of deposition/gathering have not been represented for more clarity. Panel (b) shows the domain decomposition D_2 that would be required to load balance the FFTs in the Maxwell solver step efficiently. One can see that in that case, the limits of MPI subdomains in D_1 do not coincide with the limits of MPI subdomains in D_2 due to the presence of Guardcells appended at the MPI group boundaries and needed in the computation of FFTs.

Load balancing all steps of the PIC loop thus requires keeping two different grids :

1. The first grid G_1 for particles and particles/mesh computations
2. The second grid G_2 for Maxwell solver computations.

Before the PIC loop, each MPI task is assigned to two subdomains: the first one is pertained to G_1 grid according to D_1 decomposition (associated with particles and particles/mesh computations) and the second one is pertained to G_2 according to D_2 (associated with the hybrid Maxwell solver). A data exchange protocol is determined according to the two grids overlaps: Each MPI task is "aware" of data communications patterns required to copy the data back and forth between G_1 and G_2 . This is achieved by computing:

- The intersection between the subdomains in G_1 (without guardcells) and the subdomains in G_2 (with groups guardcells).
- The intersection between the subdomains in G_2 (without guardcells) and the subdomains in G_1 (without guardcells).

Then, at each PIC iteration:

1. Fields arrays are copied from grid G_1 to grid G_2 including groups guardcells of G_2 . Overlapping grid regions pertaining to the same MPI domain on G_1 and G_2 are simply copied, while other regions of G_2 are updated using MPI communications.
2. Maxwell's equations are solved on G_2 grid.
3. Fields arrays are copied from grid G_2 to grid G_1 not including guardcells. Overlapping grid regions pertaining to the same MPI domain on G_2 and G_1 are simply copied, while other regions of G_1 are updated using MPI communications.
4. Finally, guardcells required for particles-mesh computations are exchanged between neighboring subdomains on D_1 .

The additional load balancing step, specific to the hybrid strategy, adds an extra cost to the hybrid parallelization. In practice, the data exchange step is rather computationally cheap and only results in a small overhead to the PIC loop as the majority of data exchanges are simple data copies within the same MPI process, while the actual communications exchanges between different MPIs are done through non blocking MPI send/receive calls in an efficient manner. As shown by the results of the scalability benchmarks of the hybrid solver, this time overhead is largely compensated by the gain in terms of regular MPI data exchanges previously discussed.

V.4.4 Benchmarks of the hybrid solver on different architectures

The new solver has been benchmarked on both THETA and MIRA at very large scales. The goal of these tests is to assess the strong scaling of the hybrid solver and compare it to the scaling of the local solver. All the benchmarks considered the case of a 3D homogeneous plasma with 1 pseudo-particle per cell for both ions and electrons.

Although the number of particles per cell may seem very low, it realistically describes the computational load of 3D plasma mirrors simulations where high spatial resolutions are required but with only a few particles per cell.

These benchmarks show that the hybrid parallelization strategy performs better than the local approach for large guard regions, and offers an equivalent performance for moderate guard regions. Moreover, the hybrid solver does not suffer any significant scaling efficiency loss, while the local solver showed an important decrease in scaling even for moderate guard regions. Both the pencil and slab decomposition techniques outperformed the local strategy, while the best performances were obtained with P3DFFT. Finally, the hybrid solver demonstrated an important memory gain compared to the local solver, attaining a factor of $\times 8$ when using the pencil decomposition.

Strong scaling on the MIRA cluster (BlueGene/Q)

On Mira the strong scaling of the hybrid solver was benchmarked using the pencil decomposition as FFTW-MPI did not show great scaling performances (cf figure V.11). This test involved up to 32768 IBM BlueGENE/Q nodes. Further benchmark parameters are detailed below:

FFT Library	n_x	n_y	n_z	n_{threads}	$n_{\text{proc_min}}$	$n_{\text{proc_max}}$	$n_{\text{mpi_per_group}}$
P3DFFT	256	2048	2048	4	2^{13}	2^{19}	256

One MPI task along x-axis (the non distributed axis in the FFT) is used in all runs. This benchmark (cf fig V.11) shows a very good strong scaling of the hybrid solver and a very good parallel efficiency at very large scales while the local solver suffers efficiency loss due to the increase in the total memory occupied by the fields arrays. Moreover the hybrid solver performs particularly better than the local solver for larger guardcells (central and right columns in V.11), attaining a factor of $\times 3$. The memory gain brought by the hybrid solver is also very important even for small guard regions and attains a factor of $\times 8$. Additionally, one can note that the hybrid solver involves less MPI exchanges than the local solver. Indeed, the overhead of the two grids coupling is largely compensated by the reduction in the regular MPI data exchanges costs.




Figure V.11: Local vs Hybrid scaling on MIRA with pencil decomposition. Each column represents the performances of a different number of guardcells. The blue curves represent the performances of the hybrid solver while the red ones represent the performances of the local solver. The upper row displays the memory gain brought by the hybrid solver compared to the local solver. The central row represents the total MPI exchanges costs involved in the Maxwell solver (Fields exchanges plus the two grids coupling for the hybrid solver). The lower row represents the strong scaling of the PIC loop. The missing points in the panel pertaining to the 32 guardcells case correspond to cases where the local domain sizes in each MPI task are smaller than the number of guardcells.

Strong scaling on the THETA cluster at ALCF (Intel KNL)

On THETA the strong scaling of the hybrid solver was benchmarked using both the pencil and the slab decompositions. The PICSAR library as well as P3DFFT have been compiled using the FFTW-MKL wrapper which allows the use of FFT-MKL instead of FFTW. In the tests performed on THETA, using the MKL wrapper allows for $\sim 25\%$ speed-up in FFT computations compared to pure FFTW. This benchmark has been performed with the "optimized" FFT setup where the optional reordering transposition is not performed, which allowed for an additional $\sim 20\%$ speed-up. In general, we have noted that the hybrid solver performs better with smaller groups on THETA. This is due to

the rapid performance loss of the distributed-memory FFT computations on this type of architectures. Nevertheless, the memory and computation time gains of the hybrid solver on THETA cluster were significant.

Slab decomposition strong scaling:

In this case, the simulation box was elongated along z-direction in order to support a large number of MPI tasks along this axis, as the slab decomposition allows grouping MPI tasks only along z-axis. The benchmark parameters are detailed below:

FFT Library	n_x	n_y	n_z	n_{threads}	$n_{\text{proc_min}}$	$n_{\text{proc_max}}$	$n_{\text{mpi_per_group}}$
FFT_MKL	160	160	393216	1	2^{15}	2^{18}	32




Figure V.12: Local vs Hybrid scaling on THETA with slab decomposition. Each column represents the performances of a different number of guardcells. The blue curves represent the performances of the hybrid solver while the red ones represent the performances of the local solver. The upper row displays the memory gain brought by the hybrid solver compared to the local solver. The lower row represents the strong scaling of the PIC loop. The missing points in the panel pertaining to the 32 guardcells case correspond to cases where the local domain sizes in each MPI task are smaller than the number of guardcells.

When increasing the number of processing units, only the number of MPI tasks along z-axis is increased. This setup allows to directly assess the hybrid solver scalability as adding processing units along other directions would result in a scalability behavior similar to that of the local solver. With

this setup, the hybrid parallelization performs better than the local technique, with a strong scaling efficiency of 87% between 512 and 3072 KNL nodes, while the local parallelization efficiency drops to 47% as shown by figure V.12. The maximum memory gain of the hybrid solver was $\times 2.5$, which agrees with the expected memory gain from V.7. In aggregate, the computational time of the PIC loop is decreased by a factor of 2

Pencil decomposition strong scaling:

Finally, the last benchmark performed at massively parallel scale aimed at assessing the performance of the hybrid solver with a pencil decomposition. A more important gain at both memory consumption and computational time is expected, as the pencil decomposition allows for a better mitigation of data redundancy as shown by equation V.4. The benchmark parameters were as follows:

FFT Library	n_x	n_y	n_z	n_{threads}	$n_{\text{proc_min}}$	$n_{\text{proc_max}}$	$n_{\text{mpi_per_group}}$
P3DFFT	240	6144	12288	1	2^{15}	2^{18}	64

The number of MPI tasks per group is equal to the number of MPI ranks per KNL node. This setup ensures that all MPI tasks forming each group are located at the same compute node of the machine. While this configuration is not necessarily optimal in all cases, it was sufficient to perform much better than that the local solver and to ensure a more important boost when compared to the hybrid slab case. Figure V.13 shows the strong scaling behaviors of the hybrid and the local solvers obtained from this benchmark.

This test proved that the hybrid solver ensures an excellent strong scaling in scenarios where the local solver suffers from important efficiency loss at large scales. The hybrid solver strong scaling efficiency attains at least $\sim 80\%$ even for large guard regions, while the local solver poorly scales at $\sim 30\%$ efficiency rate. At very large scales, the PIC loop is 400% faster with the hybrid solver, while memory consumption is reduced by a factor of $\sim \times 8$.




Figure V.13: Local vs Hybrid scaling on THETA with pencil decomposition. Each column represents the performances of a different number of guardcells. The blue curves represent the performances of the hybrid solver while the red ones represent the performances of the local solver. The upper row displays the memory gain brought by the hybrid solver compared to the local solver. The lower row represents the strong scaling of the PIC loop. The missing points in the panel pertaining to the 32 guardcells case correspond to cases where the local domain sizes in each MPI task are smaller than the number of guardcells.

V.5 Conclusion and prospects

In this chapter, a new parallelization strategy of the pseudo-spectral, FFT-based Maxwell solver has been presented. This technique has been fully implemented in WARP+PICSAR framework and supports periodic and absorbing boundary conditions (through Perfectly Matched Layers). It ensures an excellent strong scaling for an arbitrarily high number of guardcells and increases the maximum number of MPI processes that can be used to parallelize computations. Moreover, by reducing data redundancy, it also has huge benefits in terms of memory savings compared to the local parallelization approach for a given problem size.

The parallelization pattern exposed in this chapter can be transposed to pseudo-spectral solvers involving high order stencils and guardcells exchanges in numerical fluid dynamics, thermodynamics, or quantum physics. This new parallelization strategy has been the subject of a scientific paper in

Computer Physics Communication [85].

Chapter VI

Boosted frame as a visualization tool for UHI physics

VI.1 Context	96
VI.2 Need for an efficient visualization tool in UHI physics	96
VI.2.1 Limitations of HHG emission visualization in the laboratory frame	96
VI.3 Boosted frame visualization for PIC-UHI simulations	99
VI.3.1 Lorentz transformation in PIC codes	99
VI.4 High performance on the fly implementation of Lorentz transformation	101

VI.1 Context

With the advent of massively parallel supercomputers, multi-dimensional large scale numerical simulations are playing a growing and critical role in UHI physics achievements. Indeed, first principle numerical experiments, are not only essential to accurately identify and interpret the physical phenomena at play whenever there is no direct experimental diagnostics to probe particular observables, such as the ultra-fast plasma dynamics or the spatio-temporal electromagnetic fields profiles, but also to design future experiments by assessing optimal regimes of interactions. In this context, disposing of insightful and computationally efficient numerical diagnostics tools in PIC codes is of paramount importance to fully leverage numerical simulations.

In this chapter, we detail the development, during this thesis, of a massively parallel diagnostic tool to probe electromagnetic fields and macro-particles dynamics based on Lorentz transformation for multi-dimensional PIC simulations. The first section of this chapter highlights some limitations of standard numerical diagnostics in the context of UHI physics. Those limitations are associated with the difficulty in interpreting and analyzing data obtained with standard numerical diagnostics from PIC simulations. We then show how these limitations can be addressed by probing various numerical observables in a Lorentz boosted frame where various physical phenomena at play are more easily identified and understood.

VI.2 Need for an efficient visualization tool in UHI physics

VI.2.1 Limitations of HHG emission visualization in the laboratory frame

Driving high order harmonic generation on plasma mirrors is often performed with an obliquely incident laser pulse in order to maximize HHG efficiency [10]. Modelling the laser-plasma interaction at oblique incidence in a multi-dimensional setup poses a challenge in terms of interpreting and understanding simulation results. We illustrate this with two 2D PIC simulations of HHG on PM in two different regimes using WARP+PICSAR. The numerical/physical parameters used for both simulations is summarized in the next table VI.1:

	Physical parameters						Num params	
	Laser parameters				Plasma parameters			
	a_0	θ	FWHM	w_0	L	n_0	dx	ppcell
sim_ROM	15	55°	15fs	$4\mu m$	$\frac{\lambda_0}{8}$	$220n_c$	$\frac{\lambda_0}{285}$	36
sim_CWE	0.4	55°	15fs	$4\mu m$	$\frac{\lambda_0}{70}$	$220n_c$	$\frac{\lambda_0}{285}$	50

Table VI.1: Physical and numerical parameters for ROM (sim_ROM) and CWE (sim_CWE) harmonic generation simulations.

In the first simulation (sim_ROM), the laser intensity is sufficiently high to drive Doppler harmonic emission from PM. As the plasma density scale length L is relatively large, the plasma mirror is bent under the laser radiation pressure and acts as a focusing optics for the harmonic beam. On the other hand, in the (sim_CWE) case, harmonic emission is dominated by the CWE mechanism. In

this case, the variation of the laser intensity across the plasma surface results in different harmonic emission times along the laser-plasma interaction region, giving rise to divergent harmonic wavefronts [42]. Consequently, the spatial phases of the harmonic beams from both regimes are different (concave for CWE harmonics and convex for Doppler harmonics).

In order to observe the correlation between the plasma dynamics and the properties of the harmonic beam, one needs to take a closer look at the plasma surface during the interaction. Figure VI.1 (panels a&b) sketches one snapshot from each simulation at $T = 28T_0$. In the sim_ROM simulation, we observe that the high order harmonics are emitted by relativistic electron jets, pulled by the driving laser field in vacuum. In contrast, the CWE harmonics originate from inside the plasma medium.

In both cases, however, it is difficult to relate the harmonic generation mechanisms taking place at the plasma surface with spatial properties of high order harmonics from those pictures. In particular, we cannot directly assess the difference in terms of wavefront shapes for both ROM harmonics (focused beam) and CWE harmonics (divergent wavefronts, due to the laser intensity variation across the interaction region). This is partially due to the fact that at each instant, multiple laser cycles interact simultaneously with the plasma mirror, and it is not clear how to decouple both the spatial and temporal effects induced by the plasma mirror dynamics on the harmonic beam.




Figure VI.1: Snapshot of the harmonic emission by the plasma mirror in the laboratory frame (panels a & b) and in Bourdier frame (panels c & d). The CWE regime is sketched on the left panels and the ROM regime on the right panels;

To address this problematic, we propose to visualize the harmonic generation process in the Bourdier boosted frame where the laser is now normally incident to the plasma surface. Panels (c & d) of figure VI.1 depict the harmonic emission as well as the electrons density in the boosted frame in the CWE regime (left panel) and the ROM regime (right panel). In the relativistic case, we clearly identify the plasma mirror curvature induced by the strong laser radiation pressure. Moreover, we can clearly see that the PM curvature is directly responsible for the curvature of the wavefront of the reflected field. In the CWE simulation, we can now identify the convex harmonic wavefronts from this snapshot.

Note also that, although the laser field is normally incident to the plasma in the Bourdier frame, the Lorentz transformation naturally gives rise to a pulse-front tilt. For instance, assuming that the laser field is Gaussian, an analytical calculation of the transverse electric field expression at focus in the Bourdier frame shows that (z' is the laser transverse coordinate in the Bourdier frame) :

$$E'(x' = 0, z', t') \propto \exp\left(i\frac{\omega_0 t'}{\gamma} - \left[\frac{z' + t'c \sin \theta_0}{w_0}\right]^2 - \left[\frac{t'}{\gamma\tau}\right]^2\right) \quad (\text{VI.1})$$

where we assumed that the laser Rayleigh length is much larger than the pulse length: $Z_r \gg c\tau$ in

order to neglect the laser wavefront curvature near the laser focus. This expression shows that in the Bourdier frame, the laser pulse-front is tilted with an angle ϕ given by:

$$\phi = \arcsin(\sin \theta_0) \Leftrightarrow \phi = \theta_0$$

Physically, one can associate the presence of the PFT to the necessity for the incident laser field to irradiate the same portion of the drifting plasma in the Bourdier frame.

In this chapter, we propose a massively parallel data reconstruction strategy to retrieve multi-dimensional observables in the Bourdier frame. Therefore, while the simulation is performed in the 'laboratory', numerical diagnostics (particles trackers, fields snapshots ...) are captured in the Bourdier frame, where data analysis and visualization is more insightful regarding the physical processes at play.

More generally, our technique can be employed in any framework where the data analysis is carried out in a Lorentz boosted frame. In particular, it can also be employed in the context of laser-plasma wakefield acceleration simulations relying on the Vay Boosted frame technique [86, 87] in order to retrieve data into the laboratory frame, while the simulation is performed in Lorentz boosted frame. This physical case is inverse from the PM case but relies on the same transformation. All the parallel developments made here can thus also benefit to speed-up diagnostics of laser wakefield acceleration simulations in the boosted frame.

VI.3 Boosted frame visualization for PIC-UHI simulations

In the following, we briefly recall the Lorentz transformation on simulation quantities (i.e electromagnetic fields, sources, and macro-particles properties), then we propose a massively parallel implementation strategy in PIC codes.

VI.3.1 Lorentz transformation in PIC codes

We consider the same setup used in section III.4.1 to introduce the Bourdier boosted frame. The simulation box in the laboratory frame is given by:

$$S = \{(x, z) \in \mathbb{R}^2 \mid x \in [x_{min}, x_{max}], z \in [z_{min}, z_{max}]\}$$

The plasma occupies the domain $x_{min} < x < 0$ and the laser angle of incidence is θ_0 . The longitudinal coordinate along the plasma surface is denoted z . The Bourdier frame drifts with respect to the laboratory with a velocity $\vec{v} = c \sin(\theta_0) \vec{e}_z$ (see figure III.6). The Lorentz transformation parameters are $\beta = \frac{\|\vec{v}\|}{c} = \sin(\theta_0)$ and $\gamma = \frac{1}{\sqrt{1-\beta^2}} = \frac{1}{\cos(\theta_0)}$. In the Bourdier frame, the spatio-temporal coordinates (x', z', t') read:

$$\begin{aligned} x' &= x \\ z' &= \gamma(z - c\beta t) \\ t' &= \gamma(t - \beta \frac{z}{c}) \end{aligned} \tag{VI.2}$$

The antecedents in the Bourdier frame of the simulation box boundaries ($x, z = z_{max}$ or $z = z_{min}$) are the set of points (z', t') verifying:

$$\begin{aligned}
 z_{min} &= \gamma(z'_{min} + c\beta t') \\
 \Leftrightarrow z'_{min} &= \frac{z_{min}}{\gamma} - c\beta t' \\
 z_{max} &= \gamma(z'_{max} + c\beta t') \\
 \Leftrightarrow z'_{max} &= \frac{z_{max}}{\gamma} - c\beta t'
 \end{aligned} \tag{VI.3}$$

Therefore, in the Lorentz boosted frame, the simulation box layers are drifting with a velocity $-c\beta\vec{e}_z$. At a fixed time t' in the Bourdier frame, the numerical simulation is taking place in a domain comprised in $\{(x, z') \in [x_{min}, x_{max}] \times [\frac{z_{min}}{\gamma} - c\beta t', \frac{z_{max}}{\gamma} - c\beta t']\}$. We therefore can define a moving window in the Bourdier frame, drifting at the same velocity $-c\beta\vec{e}_z$ to keep track of the domain where the simulation in the Laboratory frame is taking place. The coordinates conversion associated to this transformation reads:

$$z'' = z' + c\beta t' \tag{VI.4}$$

Using equations VI.2, we can show that:

$$z'' = \frac{z}{\gamma} \tag{VI.5}$$

The simulation domain size along the transformed direction z is shrunk in the Bourdier frame by a factor of γ . Meanwhile, due to the time dilatation associated with the Lorentz transformation, the simulation duration is increased by a factor of γ . For instance, keeping track of a fixed simulation point (X, Z) in the laboratory frame during the whole simulation $[0, T_{max}]$ consists in tracking the image of this point given by $(X', Z'') = (X, Z/\gamma)$ over a total duration in the boosted frame given by $\delta T' = |\gamma T_{max} - \gamma \times 0| = \gamma T_{max}$.

On the other hand, the electromagnetic fields transformation writes:

$$\begin{aligned}
 E'_z(x', z'', t') &= E_z(x', z'', t') \\
 E'_x(x', z'', t') &= \gamma(E_x(x', z'', t') - c\beta B_y(x', z'', t')) \\
 E'_y(x', z'', t') &= \gamma(E_y(x', z'', t') + c\beta B_x(x', z'', t')) \\
 B'_z(x', z'', t') &= B_z(x', z'', t') \\
 B'_x(x', z'', t') &= \gamma(B_x(x', z'', t') - \frac{\beta}{c}E_y(x', z'', t')) \\
 B'_y(x', z'', t') &= \gamma(B_y(x', z'', t') + \frac{\beta}{c}E_x(x', z'', t'))
 \end{aligned} \tag{VI.6}$$

The current and charge density in the Bourdier frame are related to their laboratory counterparts

by the following transformation:

$$\begin{aligned}
 \rho'(x', z'', t') &= \gamma \left(\rho(x', z'', t') - \frac{\beta}{c} J_z(x', z'', t') \right) \\
 J'_x(x', z'', t') &= J_x(x', z'', t') \\
 J'_y(x', z'', t') &= J_y(x', z'', t') \\
 J'_z(x', z'', t') &= \gamma \left(J_z(x', z'', t') - c\beta\rho(x', z'', t') \right)
 \end{aligned} \tag{VI.7}$$

We also recall the macro-particles momentums and Lorentz factors expressions in the boosted frame:

$$\begin{aligned}
 p'_x(t') &= p_x(t') \\
 p'_y(t') &= p_y(t') \\
 p'_z(t') &= \gamma \left(p_z(t') - \beta\gamma m_p c (\gamma_p(t') - 1) \right) \\
 \gamma'_p(t') &= \gamma \left(\gamma_p(t') - 1 \right) + 1 - \frac{\gamma\beta}{m_p c} p_z(t')
 \end{aligned} \tag{VI.8}$$

where γ_p and γ'_p denote the particle Lorentz factors in the laboratory frame and the Bourdier frame respectively, and m_p the particle mass at rest.

A straightforward approach to snapshot the simulation observables in the Bourdier frame may consist in dumping large amounts of data from the simulation in the laboratory frame and then post-processing the results following equations VI.6, VI.7 and VI.8 to reconstruct the required data for each snapshot.

However, this approach is very computationally demanding both in terms of computation time and data storage. Moreover, a large portion of the dumped data in the laboratory frame would not be useful for the reconstruction step in the boosted frame.

To address those issues, we propose an efficient and massively parallel, on the fly implementation of the Lorentz boosted frames diagnostics. The next section first highlights the subtleties underlying the snapshots reconstructions in the boosted frame and details the massively parallel strategy to parallelize these computations.

VI.4 High performance on the fly implementation of Lorentz transformation

Reconstructing different data snapshots in the boosted frame on the fly requires an efficient filling strategy to keep track of the time and space couplings resulting from the Lorentz transformation.

This is illustrated in figure VI.2. The grids colored in pink represent different snapshots in the lab frame, while the ones colored in cyan represent snapshots in the boosted frame. The color opacities vary with respect to the transformed axis z (or z'') and denote the time evolution in the dual space in

each case.

At each time step in the laboratory frame, each slice of data orthogonal to the z -axis corresponds to different times in the boosted frame. Therefore, reconstructing a single snapshot on the fly requires to compute the intersection line (if it exists) between the simulation domain and the snapshot domain in the boosted frame at each time step. As the time evolves during the simulation, this intersection line drifts to larger values of z (see black and blue dashed arrows linking the lab and the boosted frames in the figure) until it covers all the simulation box in the lab frame.

The time interval (in the lab frame), during which the space time domain associated to a snapshot in the boosted frame, defined at time $T' = j\Delta T'$ intersects the simulation domain is given by:

$$t_{lab} \in \left[\frac{j\Delta T'}{\gamma} + \frac{\beta}{c}z_{min}, \frac{j\Delta T'}{\gamma} + \frac{\beta}{c}z_{max} \right] \quad (\text{VI.9})$$

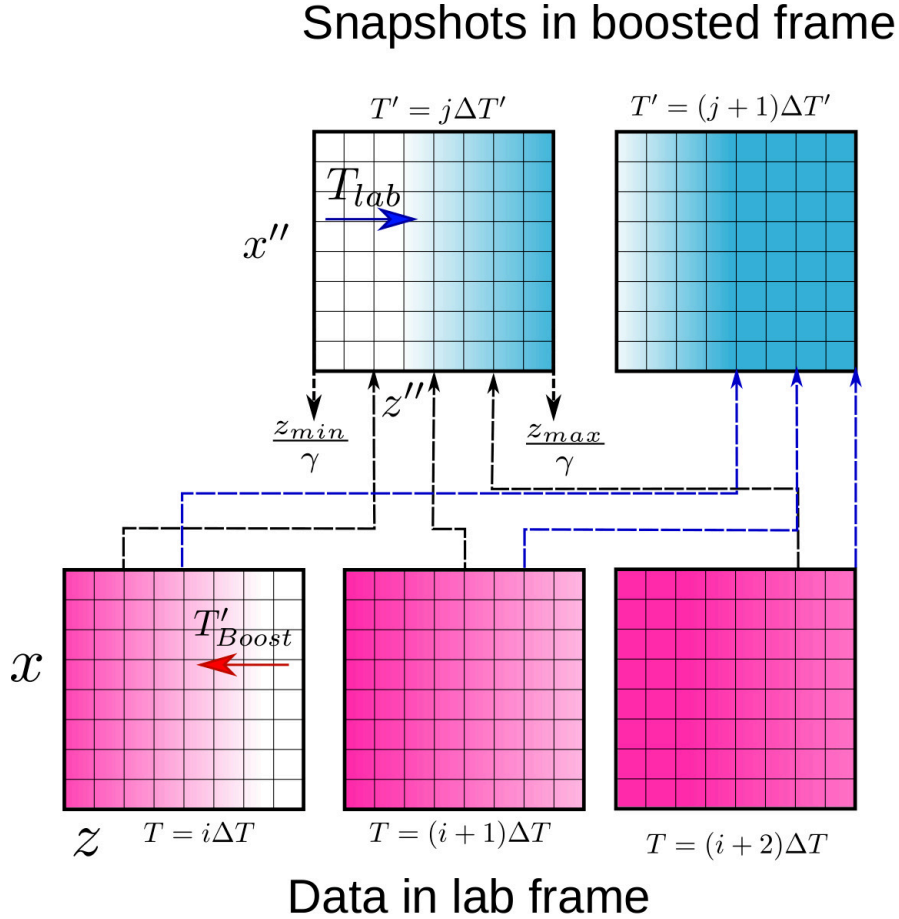


Figure VI.2: Schematic representation of the relationships between data in the boosted frame and data in the lab frame.

Taking into account the intersection slices between the laboratory and the boosted frame domains, one can perform an on the fly reconstruction of simulation observables in the boosted frame. For mesh grid quantities, the following algorithm can be used:

Algorithm 1 Boosted frame reconstruction

- 1: **for** all PIC time steps $i\Delta t$ **do**
 - 2: Perform PIC computations.
 - 3: **for** all boosted frame snapshots $j\Delta T'$ **do**
 - 4: **if** simulation time $i\Delta t$ is between lab times boundaries of the j^{th} snapshot (equation VI.9)
 then
 - 5: Compute the intersection slice between the laboratory and boosted domains.
 - 6: Extract field data along the intersection slice.
 - 7: Perform the Lorentz transformation along the extracted data.
 - 8: Write the transformed data in the file associated to j^{th} snapshot.
-

In a similar way, it is possible to snapshot pseudo-particles quantities in the boosted frame. This requires to compute at each time step $i\Delta T$ the proper time for each pseudo-particle p in the boosted frame. Whenever the boosted time of a pseudo-particle crosses a given boosted frame snapshot time, i.e:

$$t'_p(i\Delta T) < j\Delta T' \leq t'_p((i+1)\Delta T)$$

it is captured by the diagnostic and its characteristics $(x, z, \vec{p}, \vec{E}_p, \vec{B}_p)$ at times $i\Delta T$ and $(i+1)\Delta T$ in the Laboratory frame are Lorentz transformed using equations VI.8, VI.2 and VI.6 and smoothed to the exact boosted frame snapshot time $j\Delta T'$ before dumping.

Parallelization of the algorithm

As opposed to numerical diagnostics in the laboratory frame, massively parallel data reconstruction in the boosted frame is more intricate and requires a more elaborated strategy in order to be computationally efficient and scalable.

In this paragraph, we detail the parallelization strategy developed and implemented in the WARP+PIC SAR framework to parallelize the boosted frame diagnostics. This parallelization relies on the HDF5-MPI¹ library to save data on disk in parallel and is consistent with the domain decomposition employed in PIC codes.

During the PIC loop, each processing unit performs the Lorentz transformation on the data laying in its subdomain, delimited by $[z_{\text{min_local}}, z_{\text{max_local}}]$, and buffers the results associated to the desired snapshots in the boosted frame. Data is dumped on disk periodically during the simulation using HDF5-MPI library, preferably when the memory is full. The dumping periodicity may be adjusted depending on the snapshot sizes and the available free memory on each processing unit. This algorithm is schematically illustrated by figure VI.3 and detailed below:

¹<http://www.hdfgroup.org/HDF5/>

Algorithm 2 Massively parallel boosted frame reconstruction.

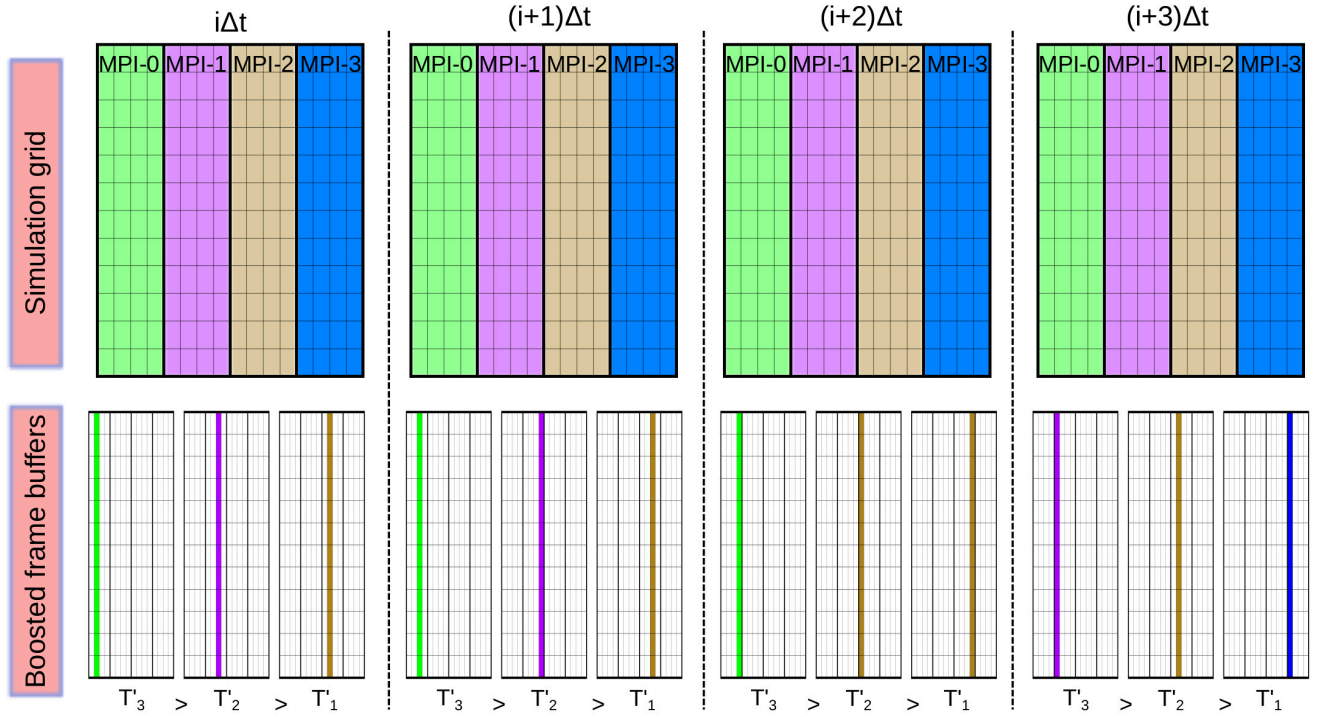
```

1: for all PIC time steps  $i\Delta t$  do
2:   Perform PIC computations.
3:   for all boosted frame snapshots  $j\Delta T'$  do
4:     if  $\frac{j\Delta T'}{\gamma} + \frac{\beta}{c}z_{\min\_local} \leq i\Delta t \leq \frac{j\Delta T'}{\gamma} + \frac{\beta}{c}z_{\max\_local}$  then
5:       Compute the intersection slice between the laboratory subdomain and boosted subdo-
        main.
6:       Extract field data along the intersection slice.
7:       Perform the Lorentz transformation along the extracted data.
8:       Save data in the buffer associated to  $j^{\text{th}}$  snapshot.
9:       if  $i\% \text{dum\_freq} = 0$  then
10:        for all boosted frame snapshots  $j\Delta T'$  do
11:          Build an MPI subcommunicator associated to the  $j^{\text{th}}$  snapshot, containing all MPI
            tasks whose  $j^{\text{th}}$  buffer is not empty.
12:          for all boosted frame snapshots  $j\Delta T'$  do
13:            if Current MPI task belongs to the  $j^{\text{th}}$  MPI subcommunicator then
14:              Open the  $j^{\text{th}}$  HDF5 file in parallel, using the corresponding MPI subcommuni-
                cator.
15:            for all boosted frame snapshots  $j\Delta T'$  do
16:              if Current MPI task belongs to the  $j^{\text{th}}$  MPI subcommunicator then
17:                Dump data in the  $j^{\text{th}}$  HDF5 file.
18:                Free my  $j^{\text{th}}$  buffer.
19:            for all boosted frame snapshots  $j\Delta T'$  do
20:              if Current MPI task belongs to the  $j^{\text{th}}$  MPI subcommunicator then
21:                Close  $j^{\text{th}}$  HDF5 file.
22:                Delete the  $j^{\text{th}}$  MPI subcommunicator.

```

Note that the loop redundancy during the write-to-disk operation (lines 12, 15, and 19 in algorithm 2) is necessary to allow for a parallel data dumping on the disk. Indeed, the loops starting in lines 12 and 19 involve synchronized MPI-based operations (Subcommunicators creations and deleting, HDF5 files opening and closing). This synchronization needs to be broken in order to perform the data dumping in parallel (loop starting at line 15). Otherwise, MPI tasks only involved in the writing of certain snapshots would be waiting for the MPI tasks involved in the writing of other snapshots to finish their job before starting to dump data in their turn.

1-Buffering step:



2-Dumping step:




Figure VI.3: Schematic principle of the massively parallel boosted frame reconstruction algorithm. During each simulation time step (buffering step), each MPI extracts, processes and buffers the data slices corresponding to different snapshots in the boosted frame (T'_3 , T'_2 , T'_1). During the dumping step, each MPI task opens the HDF5 data files associated to the snapshots for which it has buffered data during the simulation. Data is then written on disk in parallel on each file. The cells colored in gray correspond to the boosted frame domains that have already been covered and dumped into disk during the previous dumping operations.

We used this massively parallel boosted frame data reconstruction algorithm to retrieve the simulation data presented in figure VI.1 (panels c and d). For information, we recorded 400 snapshots at a rate of 20 frames per laser optical cycle (in the boosted frame). These simulations were performed on the MIRA cluster at ALCF. Each simulation involved 8192 MPI tasks distributed over 2048 com-

pute nodes. The Bourdier frame diagnostic multiplied the simulation time by a factor of ~ 3 , which constitutes a reasonable cost given the huge amount of recorded data (400 gigabytes). Note that an a-posteriori reconstruction using data from the simulation frame would have required more than 10000 snapshots and tens of Terabytes of memory storage, while increasing the computation time by many orders of magnitude.

Part C

Physics study with the new tools

Chapter VII

Generating isolated attosecond pulses in the relativistic regime

VII.1 Introduction	109
VII.2 Limitations of the attosecond lighthouse effect for Doppler harmonics	112
VII.2.1 Physical setup	112
VII.2.2 Numerical parameters	112
VII.2.3 Post-processings	114
VII.2.4 Results	115
VII.3 Effect of the plasma density scale length on the harmonic beam divergence	117
VII.4 PPM curvature compensation via laser wavefront curvature	120
VII.4.1 General principle	120
VII.4.2 Divergence of high order harmonics generated by a spatially chirped and defocused laser pulse	121
VII.4.3 Numerical validation	128
VII.4.4 Evolution of the optimal defocusing length with the gradient scale length	130
VII.4.5 Effect of the laser incidence angle	133
VII.4.6 Effect of the pulse-front tilt on the angular separation of attosecond light pulses	135
VII.4.7 Attosecond light pulses separation in 3D	141
VII.5 Laser pulse shaping	148
VII.5.1 General principle	148
VII.5.2 Flattening a laser pulse profile in the presence of spatial chirp	149
VII.5.3 PIC simulations of ROM harmonics generation with a shaped laser pulse	152
VII.6 Conclusion	155

VII.1 Introduction

When an ultra-intense ultra-short laser beam irradiates a solid or a gas target, a train of attosecond light pulses associated with a high order X-UV spectrum is emitted collinearly to the driving laser field. If properly controlled, these extremely short light bursts may have a huge impact science as they can be used for ultra-fast spectroscopy to perform a wide variety of time-resolved experiments [88].

One of the main challenges facing the development of stable and reliable attosecond light sources from plasma mirrors is the difficulty to produce single isolated pulses which are naturally more convenient for ultra-fast spectroscopy.

Many techniques have already been proposed to address this challenge, both for gas and overdense plasma high order harmonics. Multiple schemes aim at inhibiting all but one attosecond pulse in the train by confining the harmonic emission process to one laser optical cycle only. Such schemes include polarization gating and intensity gating methods.

Both of these approaches have already been successfully experimented to produce isolated attosecond pulses from gaseous media [89, 90] [91].

Nevertheless, transposing those methods to Doppler harmonics is very challenging as it would require strong control over the temporal duration, the ellipticity and the carrier-envelope phase (CEP) of an extremely intense laser beam. Therefore, producing unique attosecond light pulses from PM still remains an important issue. But thanks to the important development in laser technologies, many advances have been achieved in this field in recent years.

Intensity gating techniques take advantage of the non-linear dependency between the laser intensity and the harmonic generation efficiency. If the laser pulse is sufficiently short (one or two laser cycles), harmonics whose orders are close to the harmonic cutoff are solely generated by the optical cycle near the laser pulse maximum amplitude. As a result, the attosecond pulses train pertaining to these harmonics is dominated by a unique attosecond pulse.

In 2018, D.Kormin *et al* [92] realized the first experimental demonstration of the intensity gating technique for Doppler harmonics. This achievement was made possible by using a very short (5 fs time duration FWHM), CEP-controlled driving laser beam focused on a $\simeq 1\mu\text{m}$ focal spot to reach ultra-relativistic intensities ($a_0 = 7$). By properly controlling the laser CEP and the plasma density scale length, the authors were able to produce a train constituted of only three attosecond light spikes with an intensity contrast of at least 30.

The polarization gating method exploits the harmonic generation dependency on the driving laser ellipticity. For a circularly polarized laser pulse impinging the plasma target at normal incidence, the

laser vector potential normal to the target exhibits no fast oscillations at the laser frequency. Therefore, the ponderomotive force acting on the plasma surface follows the laser temporal envelope and does not drive any plasma surface oscillations, resulting in the absence of harmonic emission. The basic idea of the polarization gating scheme consists in appropriately inducing a time-dependent laser ellipticity, such that the high order harmonics generation process is confined to less than one laser optical cycle. This way, a single attosecond light pulse can be generated with a multi-cycle laser pulse.

Although numerous theoretical and numerical studies have been performed during the last few years [93, 94, 95], the polarization gating scheme has never been successfully implemented to produce unique attosecond X-UV light pulses from plasma mirrors in the relativistic regime (despite a recent, not fully convincing, experiment [96]).

In section III.5, we have introduced another approach, recently proposed by the LYDIL/PHI team to isolate attosecond light beams. This approach, known as the attosecond lighthouse effect, relies on introducing a spatio-temporal coupling on the driving laser field -called wavefront rotation [53]-, which induces a rotation of the laser wavefronts over time. When an ultra-intense, ultra-short laser beam, exhibiting wavefront rotation, drives a high order harmonic generation process, attosecond light pulses are emitted along the instantaneous rotating laser wave vector and angularly separate while propagating in vacuum in the far-field. The isolation of a unique attosecond pulse can be achieved simply by placing a spatial filter in the path of a single pulse (see figure III.12) [26].

This method, has already been used to produce separated attosecond light pulses from plasma mirrors in the non relativistic regime [51], and in gas media [52, 97]. These experiments exploited extremely short (FWHM $\sim 5fs$) laser beams that can sustain relatively high wavefront rotation velocities. However, the attosecond lighthouse technique has never been successfully applied to Doppler harmonics, generated on plasma mirrors in the relativistic regime.

To achieve a satisfactory angular separation of attosecond light pulses with this scheme, the rotation angle of the laser wave vector over a single optical cycle should be larger than the angular divergence of each individual attosecond light pulse. Otherwise, the attosecond pulses train do not split into isolated spikes in the far-field. Unfortunately, high order harmonics emitted in the relativistic regime are usually very divergent due to the focusing effect induced by the PM, which is curved under the non-uniform laser radiation pressure at focus [98, 99, 48]. Besides, the maximum achievable wavefront rotation velocity for a given laser beam is limited by its frequency bandwidth: for a given laser divergence, a longer pulse implies a smaller wavefront rotation velocity and thus a smaller angular separation between two successive attosecond pulses. For laser and plasma parameters optimizing HHG efficiency in the relativistic regime, the harmonic divergence can reach up to the third of the incident laser divergence. Consequently, obtaining angularly separated attosecond light pulses with the lighthouse mechanism in the relativistic regime would require extremely short laser pulses (of less than two laser cycles), which are very hard to produce and control for the extreme intensities considered here. These limitations have set a major obstacle in the experimental realization of the

lighthouse effect in the relativistic regime.

To break this barrier, we investigated feasible ways of strongly reducing harmonic beam divergence to enable the attosecond lighthouse with a standard high-power, multi-cycle laser. By mitigating the focusing effect exerted by the plasma mirror curvature on the harmonic beam in the relativistic regime, we show that it is possible to significantly reduce the harmonic divergence and achieve a clear angular separation of X-UV attosecond pulses. In this context, we have identified two realistic and feasible schemes that both lead to a good separation of attosecond pulses.

The first technique consists in curving the laser wavefronts at the plasma surface: by appropriately placing the plasma mirror out of the laser focus, we show that it is possible to produce diffraction-limited, low divergence attosecond pulses. In the presence of wavefront rotation, different attosecond pulses effectively separate in the far-field. We developed an analytical toy model that accurately predicts the best defocusing distance maximizing the separation efficiency between Doppler attosecond pulses based on the laser and the plasma parameters. The effectiveness of this technique is extensively investigated numerically using PIC simulations in 2D and 3D geometries.

The second divergence reduction technique investigated in this thesis consists in shaping the laser cross section at focus in order to directly suppress the curvature of the PM responsible for the high divergence levels of Doppler harmonics. This spatial shaping technique leverages the spatio-spectral couplings responsible for the wavefront rotation of the driving laser beam and aims at homogenizing the radiation pressure exerted by the laser on the plasma surface. By doing so, the PM is homogeneously dented over the laser-plasma interaction region, which in return, mitigates the focusing effect of the PM and limits the divergence of the attosecond pulses.

In the following, we will first review the current limitations of the lighthouse effect for Doppler harmonics and show how the PM denting is severely limiting the angular separation of attosecond light pulses. Then we detail how we can overcome these limitations and obtain isolated attosecond pulses by using novel techniques to reduce harmonic divergence.

VII.2 Limitations of the attosecond lighthouse effect for Doppler harmonics

In this section, we numerically investigate the limitations of the attosecond lighthouse effect to angularly separate Doppler harmonics.

We first show that with a physical setup that optimizes the HHG efficiency in the relativistic regime, the lighthouse scheme completely fails in angularly separating attosecond light pulses due to the important harmonic divergence increase resulting from the PM curvature. Second, we show that controlling the PM curvature by decreasing the plasma density scale length L in order to mitigate the harmonic divergence is not a good strategy to achieve a clear angular separation of attosecond light pulses as it significantly hinders the Doppler HHG efficiency.

VII.2.1 Physical setup

For this numerical experiment, we choose a physical setup that can be achieved with existing or forthcoming Petawatt laser systems, like the Apollon laser in Saclay [100] or ELI-ALPS in Hungary [101]. The laser pulse duration is 16 fs FWHM in intensity. The beam waist before introducing spatio-temporal couplings is $w_0 = 3.2\mu\text{m}$. The laser is p-polarized and has a normalized amplitude of $a_0 = 30$ (in the absence of PFT), it impinges the target with a 45° incidence angle. The pulse-front tilt amount ξ before the focusing optics maximizes the laser wavefront rotation velocity at focus (i.e. $\xi = \xi_{max} = \frac{r_0}{R}$). In this case, the laser wavefront rotation velocity reaches $V_r = 2.926\text{mrad}\cdot\text{fs}^{-1}$, and the laser intensity is decreased by a factor of 2 (and the normalized laser amplitude by a factor of $\sqrt{2}$). The plasma target is composed of a plasma bulk with a density of $n_0 = 220n_c$ and a preplasma with an exponential density ramp $n(x) \propto \exp(x/L)$ where $L = \frac{\lambda_0}{8}$. This plasma density scale length maximizes the HHG efficiency in the relativistic regime for a 45° laser incidence angle [37]. The ions have a mass of $m_i = 12 \times 1836m_e = 22,032m_e$ (with m_e the electron mass). This setup emulates an overdense plasma resulting from the full ionization of a carbon plastic target by the laser prepulse.

VII.2.2 Numerical parameters

Numerically modelling HHG on plasma mirrors with Particle-In-Cell codes in a full 3D geometry is extremely computationally-intensive as it would require tens of billions of cells/pseudo-particles. To reduce this computational cost, we first rely on PIC simulations in a 2D geometry: the simulation domain corresponds to the plane of incidence of a p-polarized laser pulse (i.e. the magnetic field is perpendicular to the simulation plane). In the presence of wavefront rotation, the laser wave vector is rotating in the simulation plane. This type of simulation still requires the use of a very fine grid sampling and a robust Maxwell equations solving algorithm. In this context, employing the ultra-high order Pseudo-Spectral Analytical Time Domain (PSATD) Maxwell solver is appropriate as it ensures a dispersion-less numerical Maxwell's equations solving up over the whole Fourier spectrum.

Unless explicitly mentioned, all numerical results presented in this chapter use the PSATD Maxwell solver with a stencil order of $n_{order} = 100$ and 8 guard cells. For these parameters, we checked

beforehand that the level of stencil truncations is sufficiently low and does not affect the physics at play [70]. The cell size is $dx = dz = \frac{\lambda_0}{285}$ and the time step is $dt = \frac{dx}{c}$. Besides, we employ the Esirkepov -charge conserving- current deposition algorithm [61] with 3rd order B splines shape functions for both current/charge deposition and field gathering. 6 particles per cell are used for both electrons and ions. Table VII.1 summarizes all the numerical and physical parameters used for this first simulation:

Physical parameters							Numerical parameters	
Laser parameters				Plasma parameters				
$a_0 \Leftrightarrow I$	θ_0	FWHM	w_0	ξ	L	n_0	dx	ppc
$30 \Leftrightarrow 1.9 \times 10^{21} Wcm^{-2}$	45°	$16fs$	$3.2\mu m$	ξ_{max}	$\lambda_0/8$	$220n_c$	$\lambda_0/285$	6

Table VII.1: Numerical and physical setup for the lighthouse+ROM simulation.

The computational cost of each 2D simulation is around 10000 core hours on a KNL-based super-computer.

A streaking probe is located $15\lambda_0$ away from the plasma surface and captures at each time step the reflected transverse magnetic field that crosses the streaking plane (see figure VII.1).

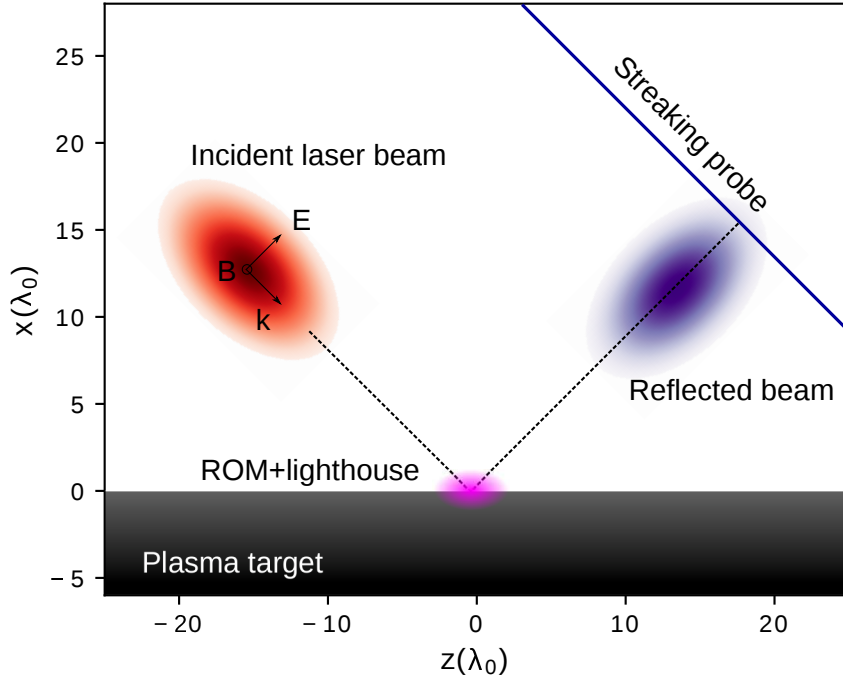


Figure VII.1: Schematic representation of a 2D attosecond lighthouse simulation. The incident laser field is p-polarized, and the wavefront rotation occurs in the simulation plane (x-z). The streaking probe is used to record the reflected electromagnetic field spatio-temporal profile.

VII.2.3 Post-processings

From the reflected field recorded by the streaking probe, we can numerically compute the angularly resolved harmonic spectrum $I_h(\theta, \omega)$:

$$\begin{aligned} I_h(\theta_x, \omega) &= |\widehat{E}(k_x, \omega)|^2 \\ k_x &= \arcsin(\theta_x) \times \frac{\omega}{c} \end{aligned} \quad (\text{VII.1})$$

This spectrum can be experimentally measured on a microchannel plate detector (MCP) after spectrally diffracting the harmonic beam on a diffraction grating. The angularly resolved harmonic spectrum gives access to the spatio-spectral properties of the reflected beam at far-field and is essential to investigate specific signatures of attosecond light pulse angular separation. In particular, it can be used to perform a Fraunhofer diffraction on the attosecond pulses train associated to an arbitrary harmonic range $[nh_{min}, nh_{max}]$. This operation is simply achieved by spectrally integrating $I_h(\theta, \omega)$ between $\omega_{min} = nh_{min}\omega_0$ and $\omega_{max} = nh_{max}\omega_0$.

$$I(\theta, nh_{min}, nh_{max}) = \int_{\omega_{min}}^{\omega_{max}} I_h(\theta, \omega) d\omega$$

This profile corresponds to the spatial profile of the attosecond pulses train in the far-field. It can be experimentally measured by imaging the spectrally integrated X-UV radiations. If a good angular separation between attosecond pulses is achieved, the far-field spatial profile $I(\theta, \omega_{min}, \omega_{max})$ is constituted of a number of angularly spaced local maximums. In this case, both the angular separation ($\Delta\theta$) between subsequent pulses, and the divergence of each individual pulses θ_n , can be retrieved from this angular profile as shown by panel (a) of figure VII.2. The divergence of a single attosecond light pulses is defined as the half width at $1/e^2$ of an isolated spike and can be numerically calculated by using a Gaussian fit method for example. Otherwise, if the attosecond pulses are not sufficiently separated, then the harmonic radiation spans over a continuous angular range pertaining to the rotating laser beam numerical aperture (cf figure panel b): in this case the angular separation and the harmonic divergence cannot be deduced.




Figure VII.2: Far-field angular profiles constituted of three spikes. Panel (a): angularly separated case. Panel (b): non separated case.

Moreover, numerical simulations enable additional diagnostics that are hardly accessible in real experiments. For example, the spatio-temporal profile of the attosecond pulses train is obtained by filtering a specific harmonic range from the reflected beam, captured by the streaking probe. This train can then be propagated at any arbitrary distance from the plasma target (to directly observe the attosecond pulses separation at far-field) or to the plasma mirror emission plane (to assess the attosecond light pulses spatial profiles following the HHG emission process), using the plane-wave decomposition method [102]:

$$B_y(x_f, t, z = \Delta z) = \text{FFT}_{k_x, \omega}^{-1} \left[\text{FFT}_{x_f, t} [B_y(x_f, t, z = 0)] \times \exp \left(i \frac{\omega}{c} \Delta z - i \Delta z \sqrt{(\omega/c)^2 - k_x^2} \right) \right] \quad (\text{VII.2})$$

VII.2.4 Results

Results of this first PIC simulation (cf table VII.1) are represented on figure VII.3 and VII.4. [Panel (a) Figure VII.3] illustrates the spatio-temporal profile of the harmonic beam between orders 20 and 30, propagated to a distance of $D = 100Z_r \simeq 4000\mu m$. As one can notice, the attosecond pulses train does not split at far-field. Thus, it would not be possible to select a unique attosecond light pulse. Panel (b) represents the angular profile of this harmonic train. The harmonic emission spans over a continuous angular range. The absence of clear, isolated local maximum in this profile emphasizes the absence of angular separation of attosecond light pulses. The black dashed line highlights the emission angle of the X-UV intensity peak. One can notice that the maximum harmonic emission does not occur along the specular direction $\theta = 0$. Instead, it takes place along the direction given by $\theta_{max} = 22$ mrad. To explain this observation, we represent in panel (c) the temporal profiles of the attosecond pulses train (blue curve) along with the temporal envelope of the reflected laser (red dashed curve). From these temporal profiles, one can notice that the maximum laser amplitude does not coincide with the maximum X-UV emission amplitude. This means that the HHG efficiency over time drops even though the laser intensity continues to grow. This is the signature of the relativistic gyromagnetic effect [103]. For ultra-relativistic laser intensities, the laser magnetic field strongly bends the plasma surface electrons trajectories when they are pulled back by the laser field, reducing the velocity of the plasma mirror surface oscillations responsible for the Doppler radiations, which results in a drop of HHG efficiency. This effect grows in time as the laser amplitude increases and the magnetic force becomes stronger. Therefore, the most intense attosecond pulse is usually not generated by the most intense laser optical cycle in the ultra-relativistic regime. Thanks to the wavefront rotation mechanism, the temporal and angular harmonic profiles are mutually linked. In our case, the harmonic emission amplitude maximum is reached three laser periods before the laser intensity peak, as shown by panel (c). This temporal delay corresponds to an angular shift of $\delta\theta = 3 \times T_0 \times V_r \simeq 23$ mrad (T_0 is the laser period), which agrees well with the emission angle of the X-UV intensity peak θ_{max} .




Figure VII.3: Panel (a): Attosecond pulses train propagated at a distance f ($D = 100Z_r \simeq 4000\mu m$) of the plasma target. Panel(b): angular profile of the attosecond pulse train. Panel (c): Temporal profiles of the attosecond pulses train (blue curve) and the reflected laser field (red dashed curve) taken at the plasma surface.

Panel (a) of figure VII.4 represents 3 attosecond pulses retro-propagated to the plasma surface. The attosecond pulses wavefronts are strongly curved at the harmonic emission plane. This curvature results from the PM denting induced by the inhomogeneous laser radiation pressure exerted on the plasma surface [48, 99]. In these conditions, the emitted harmonics are focused by the curved PM surface. This effect is associated with an increase in the harmonic divergence, which prevents a clear angular separation of attosecond light pulses with the lighthouse effect.

It may be interesting to notice that the plasma mirror curvature is also encoded in the spatial profile of the attosecond pulses train in far-field. Panel (b) of the same figure sketches the angularly resolved harmonic spectrum for harmonic orders between 20 and 30. One can note that the harmonic rays are tilted in the (θ, ω) space. This tilt is the signature of angular dispersion: different colors associated to a single harmonic ray propagate along different directions. This spatio-temporal coupling cannot solely result from the laser field spatio-temporal couplings. Indeed, as the laser beam is focused at the plasma surface, its spatio-temporal coupling is a pure spatial chirp. However, as the plasma surface

is strongly curved, the harmonic wavefronts are simultaneously rotating (due to the laser WFR) and curved (due to the PM curvature). As discussed in the introductory chapter (paragraph III.5.3) these two effects result in the angular dispersion of each high order harmonic (i.e $c_{k\omega} = \frac{1}{4} \frac{a_{xt}}{a_{xx}a_{tt} + a_{xt}^2}$)




Figure VII.4: Panel (a): Temporal envelopes for of three attosecond light pulses pertaining the harmonic orders between 20 and 30 at the plasma surface. Panel(b): Angularly resolved harmonic spectrum for harmonic orders between 20 and 30.

VII.3 Effect of the plasma density scale length on the harmonic beam divergence

The numerical experiment presented in the previous section shows that under realistic conditions, Doppler harmonics are too divergent to be angularly separated by the lighthouse effect. This high divergence results from the focusing effect exerted by the PM on the harmonic beam. To enable attosecond pulses separation with the lighthouse scheme, a good control over the harmonic beam divergence is thus required.

Such control can be achieved by tuning the plasma density scale length in experiments. Indeed, for a steeper plasma density scale length, the laser field penetrating the PM surface interacts with denser plasma regions. The recalling electrostatic force resulting from the ions-electrons separation, which is proportional to the plasma density, is increased. Therefore, the maximum excursion of the plasma surface electrons during the pushing phase is shorter, which results in a shallower plasma mirror denting. Consequently, the PM curvature is flatter, and the focusing effect on the high order harmonics is mitigated. Hence, for shorter plasma density scale lengths, the attosecond pulses are less divergence.

Nevertheless, this divergence decrease comes with a drop of the HHG spectrum roll-off and cut-off [104]. Indeed, as the restoring force increases for sharper density gradients, the outgoing electrons of the PM surface responsible for the Doppler harmonic emission gain less energy from the driving laser

field, which decreases the HHG efficiency.

In this section, we numerically investigate the relationship between the plasma scale length L the angular separation of attosecond pulses with the lighthouse mechanism. To achieve this, we performed a series of 2D PIC simulations where we varied the plasma scale length L between $\frac{\lambda_0}{15}$ and $\frac{\lambda_0}{200}$. Other numerical and physical parameters are the same as the ones of table VII.1.

Results from these simulations are represented in figure VII.5, each column corresponds to one plasma density scale length case (left $L = \frac{\lambda_0}{15}$, center $L = \frac{\lambda_0}{40}$, right $L = \frac{\lambda_0}{200}$). Panels (a,b,c) illustrate the angularly resolved harmonic spectra for each case. As one can note, the HHG cut-off is strongly affected by the plasma density scale length decrease. Panels (d,e,f) represent the angular profiles for harmonic orders between 17 and 33. The attosecond pulses trains associated to these harmonics are sketched in panels (g,h,i). Decreasing the plasma scale length results in the appearance of subsequent local maximums in the angular profiles of the high order harmonic radiations (for $L = \frac{\lambda_0}{40}$ and $L = \frac{\lambda_0}{200}$). These angular profiles are associated to trains of successive, low divergence attosecond light pulses (panels (h) and (i)), as opposed to the attosecond pulses train obtained from a relatively long plasma density scale length ($L = \frac{\lambda_0}{15}$), panel (g) which is composed of high divergent attosecond light pulses. However, the decrease in the harmonic divergence is accompanied with a strong drop in the HHG efficiency, as shown by the amplitude decrease of the angular profiles. Finally, by computing both the harmonic divergence and the attosecond light pulses angular separation, we find that the separation ratios $\frac{\Delta\theta}{\theta_n}$ are 0.9 and 0.78 for $L = \frac{\lambda_0}{40}$ and $L = \frac{\lambda_0}{200}$ respectively, which is not sufficient to efficiently separate attosecond light pulses.

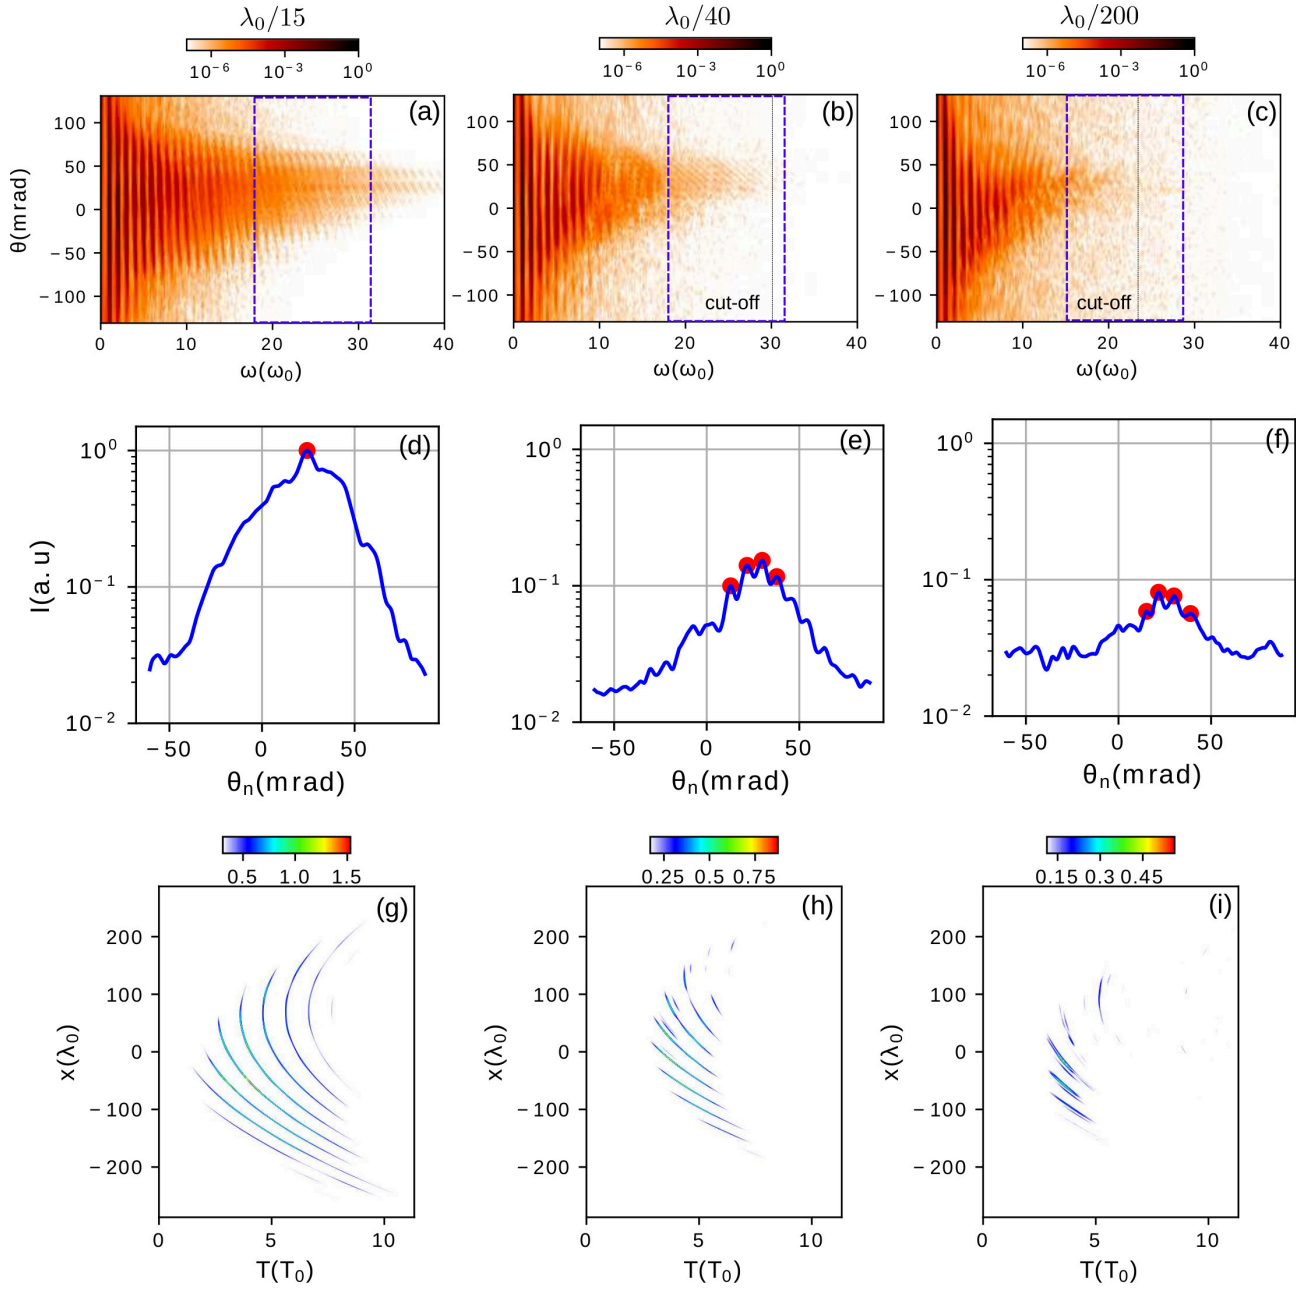


Figure VII.5: Angularly resolved harmonic spectra (first row), angular profiles for harmonic orders between 17 and 33 (second row) and spatio-temporal profiles of attosecond light pulses of harmonic orders between 17 and 33 propagated far from the plasma target ($D = 100Z_r \simeq 4000\mu m$) for three values of plasma scale lengths.

VII.4 PM curvature compensation via laser wavefront curvature

VII.4.1 General principle

In this section, we detail the first harmonic divergence reduction technique explored in this thesis. This technique is rather simple, and should be easily implemented within HHG experiments. It could lead to a substantial drop in harmonics divergence.

This method consists in placing the PM out of the laser focus. This way, the laser wavefronts impinging the PM are curved. This curvature brings an additional quadratic term ϕ_L to the harmonic spatial phases at the PM surface (c.f paragraph III.4.10). This term can be used to decrease/control the harmonic divergence, depending on which side of the laser focus the PM is placed.

If the driving laser beam is focused slightly **before** reaching the plasma surface (cf figure VII.6 panel (b)), ϕ_L is then of opposite sign than the harmonic spatial phase induced by the PM curvature ϕ_p . In this case, the usual PM focusing effect exerted on high order harmonics is mitigated by the laser wavefront curvature. Additionally, if the total harmonic quadratic phase completely vanishes (i.e. $\phi_n = \phi_L + \phi_p = 0$), then the harmonic divergence should be reduced to the diffraction-limited divergence given by $\theta^0 = \frac{\lambda_0}{n\pi w_n}$ (where n is the harmonic order and w_n is the harmonic beam size of the n^{th} harmonic order at the PM surface). Coupled with the wavefront rotation mechanism, this technique could lead to a clear angular separation of attosecond light pulses in the relativistic regime.

To assess the effectiveness of this scheme, it might be interesting to visualize the high order harmonic emission with a defocused laser beam in the Bourdier frame, using the visualization technique developed in chapter VI. For this, we performed two 2D-PIC simulations using regular Gaussian laser (i.e. without wavefront rotation). In the first simulation the laser is focused at the plasma surface (cf figure VII.6 panel (a)), while in the second one, the laser focus is located $-0.6Z_r$ before the plasma surface. This defocusing distance induces a laser wavefront curvature that matches the PM curvature. Additional physical/numerical parameters are detailed in the following table VII.2:

Physical parameters						Num params			
Laser parameters					Plasm param				
a_0	θ	FWHM	w_0	$\xi(\xi_{max})$	$\Delta z(Z_r)$	L	n_0	dx	ppcell
10	45°	16fs	3.2 μm	0	-0.6	$\lambda_0/15$	220 n_c	$\lambda_0/200$	8

Table VII.2: Physical/numerical parameters for the Bourdier visualization.

Panels (c) and (d) of figure VII.6 sketch the emission of a single attosecond light pulse associated to harmonic orders between 16 and 24 near the PM surface. In the first case, where the laser wavefronts are flat, the attosecond light pulses are curved at the PM surface level. On the other hand, when the laser is focused before the target, the wavefronts of the attosecond light pulses are flat, showing that focusing effect induced by the PM curvature has effectively been suppressed.




Figure VII.6: Attosecond pulses emission by the PM in Bourdier frame. Panels (a) and (b): schematic representation of driving laser wavefronts. Panels (c) and (d): emission of a single attosecond light pulse when the laser is focused at the PM surface and $0.6Z_r$ before the PM surface respectively.

VII.4.2 Divergence of high order harmonics generated by a spatially chirped and defocused laser pulse

In this section, we derive a theoretical model assessing the Doppler harmonic divergence θ_n as a function of the defocusing distance and additional laser and plasma parameters in the presence of wavefront rotation. Based on this model, it is then possible to compute the separation quality ratio of Doppler harmonics $\frac{\Delta\theta}{\theta_n}$ as a function of laser and plasma parameters and deduce optimal conditions favoring attosecond pulses separation. The accuracy of this model is then validated by 2D and 3D PIC simulations.

As detailed in paragraph III.4.10 and in reference [48], under the paraxial approximation, the spatial phase of Doppler harmonics emitted by a PM illuminated by an ultra-intense defocused Gaussian

laser beam is given by:

$$\phi_n(x) = \frac{2n\pi}{\lambda_0} \left[\frac{x^2}{2R(\Delta z)} + \frac{x^2}{2\cos(\theta)f_p(\Delta z)} \right] \quad (\text{VII.3})$$

where x denotes the transverse position at the PM surface, Δz the distance between the laser focus and the PM surface, $f_p(\Delta z)$ is the PM radius of curvature under the laser radiation pressure, and $R(\Delta z)$ is the laser wavefronts radius of curvature. In this case, the harmonic divergence is given by the following equation:

$$\begin{aligned} \theta_n(\Delta z) &= \theta_n^0(\Delta z) \sqrt{1 + (n\psi(\Delta z))^2 \left(1 + \frac{f_p(\Delta z)\cos(\theta)}{R(\Delta z)}\right)^2} \\ \psi(\Delta z) &= \frac{2\pi}{\cos(\theta)} \left(\frac{w_n(\Delta z)}{w_0(\Delta z)}\right)^2 \delta_p(\Delta z) \end{aligned} \quad (\text{VII.4})$$

where $\theta_n^0 = \frac{\lambda_n}{\pi w_n(\Delta z)}$ is the harmonic diffraction-free divergence, $w_0(\Delta z)$ and $w_n(\Delta z)$ are the laser and harmonic beam sizes at the PM surface respectively, ψ is the PM dimensionless focusing parameter and $\delta_p = \frac{w_L(\Delta z)}{2f_p(\Delta z)}$ is the PM denting.

Equation VII.4 shows that the harmonic divergence θ_n is increased in the presence of a harmonic quadratic spatial phase ϕ_n . If the quadratic spatial phase is null, (i.e. if $f_p(\Delta z)\cos(\theta) = -R(\Delta z)$), then the harmonic divergence is equal to the diffraction-free divergence θ_n^0 , and the PM focusing effect is completely inhibited. Note also that by defocusing the laser pulse, the laser beam size at the PM surface $w_0(\Delta z)$ is larger. Therefore, the harmonic beam sizes at the emission plane $w_n(\Delta z)$ increase too, and the harmonic diffraction-free divergence is further decreased.

The PM focal distance $f_p(\Delta z)$ varies as a function of time during the laser-plasma interaction process. However, during the first laser optical cycles of the interaction, the ions motions contribution to the PM denting can be neglected. In this case, the PM focal distance is given by:

$$f_p(\Delta z) \simeq \frac{w_0^2(\Delta z)}{4L\cos^2(\theta)} \quad (\text{VII.5})$$

For relatively long laser-plasma interaction durations (tens of laser optical cycles), the PM denting is increased, and the PM focal distance is reduced by up to a factor of two. However, unless explicitly mentioned, we use the PM focal distance given by equation VII.5 in the remaining of this thesis, as we will mainly observe attosecond pulses emitted by the PM during the first laser optical cycles of the interaction.

On the other hand, the laser radius of curvature is given by:

$$\begin{aligned} R(\Delta z) &= \Delta z \left[\Delta z + \left(\frac{Z_r}{\Delta z}\right)^2 \right] \\ Z_r &= \pi \frac{w_0^2}{\lambda_0} \end{aligned} \quad (\text{VII.6})$$

where Z_r is the Rayleigh length of the laser beam. The laser wavefront curvature (i.e, the amplitude of the laser quadratic spatial phase, proportional to $\frac{1}{R\Delta z}$) increases with respect to the laser defocusing

distance Δz . It is maximum for $\Delta z \pm Z_r$. In this study, we are only interested in **negative** defocusing distances that result in convex laser wavefronts (like the one in panel (b) of figure VII.6). For a regular Gaussian laser beam, the spatial phase compensation may only occur in the range $\Delta z \in [-Z_r, 0]$. On the other hand, the PM focusing effect, which is inversely proportional to the PM focal distance $f_p \cos(\theta)$, increases with respect to the plasma density scale length L . Moreover, efficient Doppler HHG is triggered with relatively long plasma scale lengths, that induce a strong PM curvature. Therefore, it might be interesting to calculate the maximum PM density scale length inducing the maximum PM curvature that can be compensated by the laser wavefront curvature. Such a plasma density scale length L_{max} can be computed by equating the PM focal distance with the laser radius of curvature for $\Delta z = -Z_r$.

$$\begin{aligned}
 -R(-Z_r) &= f_p(-Z_r) \cos(\theta) \\
 \Leftrightarrow 2Z_r &= \frac{2w_0^2}{4L_{max} \cos(\theta)} \\
 \Leftrightarrow \frac{2\pi w_0^2}{\lambda_0} &= \frac{2w_0^2}{4L_{max} \cos(\theta)} \\
 \Leftrightarrow L_{max} &= \frac{\lambda_0}{2\pi \cos(\theta)}
 \end{aligned} \tag{VII.7}$$

For an incidence angle of $\theta_0 = 45^\circ$, we find $L_{max} \simeq \frac{\lambda_0}{4.5}$, which means that this technique is suited for Doppler harmonics, as the plasma density scale lengths of interest do not exceed L_{max} .

Figure VII.7 shows the evolution of $\theta_n(\Delta z)$ given by equation VII.4 for the 20th harmonic order for different values of plasma density scale lengths ($L = \frac{\lambda_0}{15}$ blue curve, $L = \frac{\lambda_0}{30}$ red curve, $L = \frac{\lambda_0}{100}$ magenta curve). In each case, we assumed that the harmonic beam size at the plasma surface is $w_n = 0.8w_0(\Delta z)$ and that the laser incidence angle $\theta = 45^\circ$. As the PM curvature is increased (by increasing the plasma density scale length L), the laser curvature needs to be increased further to maintain a null harmonic spatial phase at the plasma surface. Therefore, the optimal defocusing distance is larger for larger plasma density scale lengths. Note also that the minimal harmonic divergence is relatively independent of the PM density scale length L . As the PM focusing effect is completely suppressed, the PM curvature no longer affects the harmonic divergence, which shows that this divergence reduction technique allows manipulating the harmonic divergence without affecting the HHG efficiency.




Figure VII.7: Evolution of harmonic divergence as a function of the laser defocusing length for three values of plasma gradient.

In the presence of wavefront rotation, the harmonic divergence formula expressed by equation VII.4 is valid. However, different terms involved in this formula need to be recalculated in order to take into account for the laser wavefront rotation (refer to appendix A for the full derivation of the following quantities).

In the presence a wavefront rotation associated to a pulse-front tilt amount before focus ξ , the laser waist out of focus reads:

$$w_{eff}(\Delta z, \xi) = w_{eff}(\xi) \sqrt{1 + \left(\frac{\Delta z}{Z_r^{eff}} \right)^2} \quad (\text{VII.8})$$

Z_r^{eff} is the effective Rayleigh length given by:

$$Z_r^{eff} = \frac{\pi w_{eff}^2}{\lambda_0} \quad (\text{VII.9})$$

and w_{eff} the effective laser waist at focus:

$$w_{eff}(\xi) = w_0 \sqrt{1 + \left(\frac{\xi}{\xi_{max}} \right)^2}$$

with ξ_{max} the pulse-front tilt amount that maximizes the laser WFR velocity at focus:

$$\xi_{max} = \frac{\tau L}{w_0}$$

Due to the increase in the laser beam size in the presence of a spatial chirp, the laser radius of curvature reads:

$$R(\Delta z, \xi) = \Delta z \left[1 + \left(\frac{Z_r^{eff}}{\Delta z} \right)^2 \right] \quad (\text{VII.10})$$

The maximum laser curvature is reached for a defocusing distance of $\Delta z \pm Z_r^{eff}$. Hence, in the presence of wavefront rotation, the harmonic spatial phase compensation could be achieved for defocusing distances beyond Z_r (but below Z_r^{eff}).

Moreover, it can be safely assumed, in the context of Doppler HHG, that the laser incidence angle θ is roughly constant over time. Indeed, over the total laser temporal duration, the laser wave vector total rotation does not exceed a few tens of mrad. For instance, assuming the physical parameters given by table VII.1, and supposing that the laser beam is ten optical cycles long, the laser wave vector rotates by:

$$10 \times \Delta\theta = 10T_0V_r \simeq \frac{\pi}{40} \ll \theta_0 = \frac{\pi}{4}$$

where V_r is the wavefront rotation velocity.

Finally, due to the increase in the laser beam size at the PM surface, the PM focal distance reads:

$$f_p(\Delta z, \xi) = \frac{w_{eff}(\Delta z, \xi)^2}{4L\cos^2(\theta)} \quad (\text{VII.11})$$

By injecting equations VII.8, VII.10 and VII.11 into equation VII.4, the generalized harmonic divergence $\theta_n(\Delta z, \xi)$ in the presence of wavefront rotation and a laser wavefront curvature is given by:

$$\begin{aligned} \theta_n(\Delta z, \xi) &= \theta_n^0(\Delta z, \xi) \sqrt{1 + (n\psi(\Delta z, \xi))^2 \left(1 + \frac{f_p(\Delta z, \xi)\cos(\theta)}{R(\Delta z, \xi)}\right)^2} \\ \psi(\Delta z, \xi) &= \frac{2\pi}{\cos(\theta)} \left(\frac{w_n(\Delta z, \xi)}{w_0(\Delta z, \xi)}\right)^2 \frac{\delta_p}{\lambda_0} \end{aligned} \quad (\text{VII.12})$$

On the other hand, the laser wavefront rotation velocity as a function of the defocusing distance Δz is given by the following equation (see appendix A):

$$\begin{aligned} V_r(\Delta z, \xi) &= \frac{1}{k_0} \frac{\zeta}{1 + \left(\frac{\Delta z}{Z_r^{eff}}\right)^2} \\ \zeta_f &= \frac{2\xi}{\tau_0^2 \left[1 + \left(\frac{\xi}{\xi_{max}}\right)^2\right]} \end{aligned} \quad (\text{VII.13})$$

The angular separation between two successive attosecond light pulses is given by:

$$\Delta\theta(\Delta z, \xi) = T_0 \times V_r \quad (\text{VII.14})$$

Figure VII.8 sketches the evolution of the laser wavefront rotation velocity as a function of the defocusing distance Δz for a pulse-front tilt amount that maximizes the laser wavefront rotation velocity at focus $\xi = \xi_{max}$. One can note that the laser wavefront rotation velocity loss out of focus is relatively small. For instance, at a defocusing distance of $\Delta z = \pm Z_r$, there is "only" 20% drop and $\simeq 50\%$ at $\Delta z = \pm 2Z_r$. This drop should thus not affect the effectiveness of our divergence reduction technique in improving the angular separation of attosecond light pulses by the lighthouse scheme.




Figure VII.8: Laser wavefront rotation velocity as a function of the defocusing distance Δz for $\xi = \xi_{max}$ in arbitrary units.

Using equations VII.13, VII.12 and VII.14, the attosecond pulse separation ratio $\frac{\Delta\theta(\Delta z, \xi)}{\theta_n(\Delta z, \xi)}$ can be calculated for any laser and plasma setup:

$$\frac{\Delta\theta(\Delta z, \xi)}{\theta_n(\Delta z, \xi)} = \frac{\frac{T_0 \times \zeta_f}{k_0} \frac{1}{1 + \left(\frac{\Delta z}{Z_r^{eff}}\right)^2}}{\theta_n^0(\Delta z, \xi) \sqrt{1 + (n\psi(\Delta z, \xi))^2 \left(1 + \frac{f_p(\Delta z, \xi) \cos(\theta)}{R(\Delta z, \xi)}\right)^2}} \quad (\text{VII.15})$$

It is then possible to fetch optimal laser and plasma parameters to improve the attosecond pulses separation. Note, however, that the harmonic beam sizes w_n should be fed to this model. As there is no known theoretical model to predict this quantity, we will extract it from PIC simulations and use it later in our computations. Note also that this quantity only affects the separation ratio $\frac{\Delta\theta}{\theta_n}$, but not the optimal defocusing distance.

Calculation of harmonic beam sizes

Assessing harmonic beam sizes at the PM surface is necessary in order to exploit our analytical divergence model, given by equation VII.15. To compute these quantities, we performed a 2D PIC simulation using a regular Gaussian laser field. The simulation parameters are summarized in the next table VII.3:

Physical parameters						Num params			
Laser parameters			Plasm params						
a_0	θ	FWHM	w_0	$\xi(\xi_{max})$	$\Delta z(Z_r)$	L	n_0	dx	ppcell
20	45°	16fs	3.2μm	0	0	$\lambda_0/15$	220 n_c	$\lambda_0/285$	6

Table VII.3: Physical/Simulation parameters to estimate the harmonic beam sizes.

The laser intensity at the plasma surface is roughly equivalent to that obtained with a spatially chirped laser for $\xi = \xi_{max}$ and $a_0 = 30$.

The harmonic beam sizes for orders $n \leq 25$ are sketched in figure VII.9. Due to the very large laser amplitude ($a_0 = 20$), the high order harmonics are generated on a relatively large part of the laser beam size $\frac{w_n}{w_0} \simeq 0.8$. Moreover, we note that this ratio remains constant over a relatively large harmonic range. In the remaining of this thesis, we will use this value and assume that it does not depend on the harmonic order n .




Figure VII.9: Harmonic beam sizes obtained from PIC simulation, normalized to the laser waist.

Optimal defocusing distance for attosecond pulses separation

We fix the laser incidence angle to 45° and the plasma density scale length $L = \frac{\lambda_0}{15}$. The pulse-front tilt $\xi = \xi_{max}$ maximizes the laser wavefront rotation at focus.

As validated through PIC simulations, we can reasonably assume that the harmonic beam size is $w_n = 0.8w_0$, and that it does not depend on harmonic order. The expected evolution of θ_n as well as $\Delta\theta$ is sketched in figure VII.10 for 4 different harmonics orders. Panel (a) shows that the harmonic divergences reach a minimum for $\Delta z \simeq -Z_r$. For this defocusing distance, the separation ratio $\frac{\Delta\theta}{\theta_n}$ is maximized (cf panel c). This ratio is larger than 1 for all harmonic orders beyond 10 and is even larger than 2 for $n = 25$, which suggests that an excellent separation can be reached by this technique. Furthermore, we observe that the optimal defocusing distance is independent of the harmonic order. This observation is important as it implies that minimizing the divergence of a single harmonic order is equivalent to minimizing the divergence of individual attosecond light pulses comprising multiple harmonics.

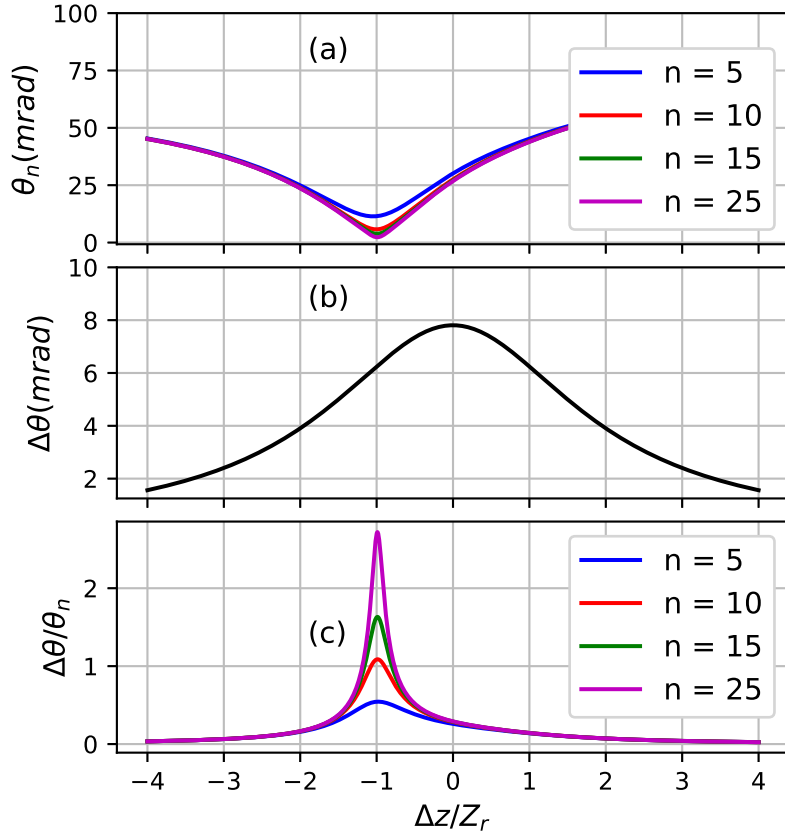


Figure VII.10: Panel (a): Divergence of 4 harmonics orders as a function of Δz . Panel (b): Angular shift between successive laser wavefronts $\Delta\theta$. Panel (c): Evolution of the ratio $\Delta\theta/\theta_n$ as a function of Δz .

VII.4.3 Numerical validation

In order to investigate the effectiveness of our harmonic divergence reduction technique in improving the angular separation of the lighthouse scheme, we carry out an extensive numerical study using 2D PIC simulations. This study is also intended to assess the validity of our analytical divergence model for a broad range of laser and plasma parameters and prepare the ground for a more expensive full 3D simulation.

Numerical validation of the laser defocusing technique in 2D

First, we perform a series of 2D PIC simulations in which we varied the laser defocusing distance Δz_r in the presence of wavefront rotation. This parameter scan is constituted of 8 simulations in which the laser defocusing distance takes the following values :

$$\frac{\Delta z}{Z_r} = \{0, -0.25, -0.5, -0.7, -0.85, -1, -1.25, -1.5\}.$$

The plasma density scale length is fixed to $L = \lambda_0/15$. Other physical and numerical parameters are the same given by table VII.1 (45° incidence angle, $a_0 = 30$).

Figure VII.11 sketches the harmonic angular profiles for harmonic orders between 20 and 30 obtained from each simulation (the defocusing distance in each case is specified in each panel). As we keep increasing the defocusing distance (in absolute value), the harmonic angular profile moves from a bell-shaped nearly concave curve (for $\Delta z = 0$ for example) to a dented shape with multiple local maximums. This shape is very obvious for the cases $-1.25Z_r \leq \frac{\Delta z}{Z_r} \leq -0.85Z_r$. It suggests that in these particular cases, the harmonic divergence θ_n has decreased significantly to a point where it dropped below the angular separation threshold $\Delta\theta$. By further increasing the laser defocusing distance, the spatial phase matching between the PM curvature and the laser wavefront curvature is no longer fulfilled, which gives rise to an increase in the harmonic divergence and a drop of in the separation ratio.

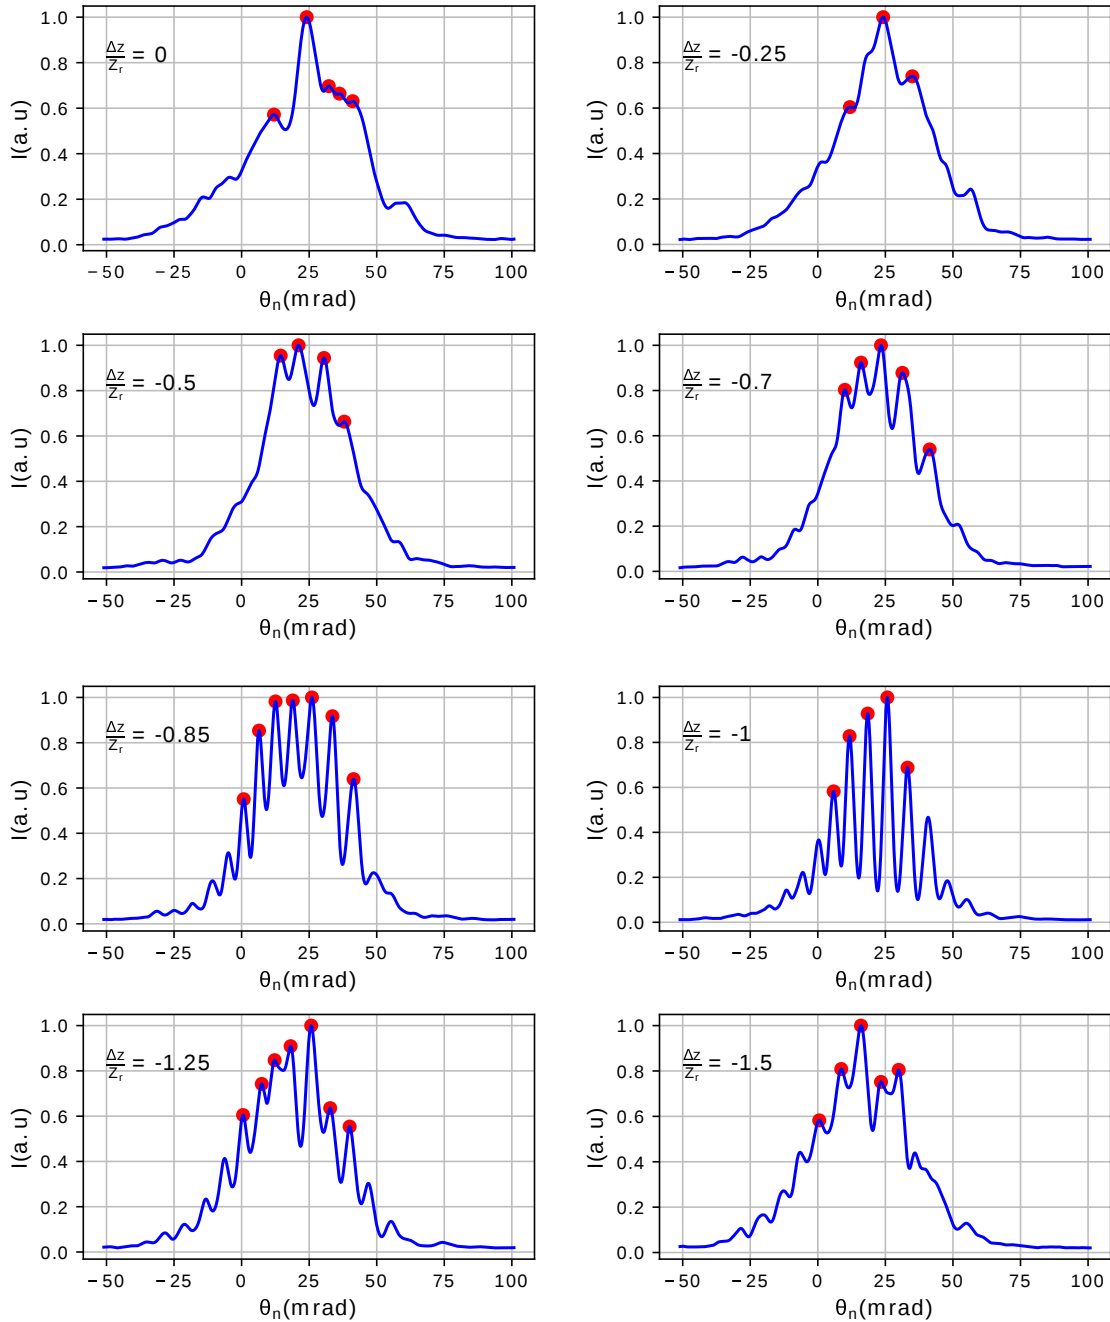


Figure VII.11: Integrated angularly resolved harmonic spectra obtained from 2D PIC simulations for different values of defocusing lengths.

In figure VII.12, the evolution of the ratio $\frac{\Delta\theta}{\theta_n}$ extracted from these PIC simulations as a function of the defocusing distance is represented in green line. In black dashed line we represent the theoretical evolution of this ratio. There is a good agreement between our analytical calculations and the PIC simulation results. The optimal separation ratio is attained for a defocusing distance of $\Delta z = -Z_r$, in agreement with the theoretical predictions (of $\Delta z = -0.99Z_r$). In this case, the obtained angular separation ratio $\frac{\Delta\theta}{\theta_n} = 2.1$ which represents a large enough contrast to produce well separated attosecond light pulses. Note finally that the gap between the theoretical and numerical separation ratios on both sides of the curve (for $\Delta z = -1.5Z_r$ and $\Delta z = -0.5Z_r$) is due to the fact that the fitting curve method used to retrieve the standard deviations of the local spikes from the harmonic angular profiles is not very robust when the separation ratio is too weak.




Figure VII.12: Separation quality as a function of the laser defocusing length. In green: the data obtained from simulations. In dashed black: the theoretical separation predicted by the toy model.

VII.4.4 Evolution of the optimal defocusing length with the gradient scale length

As discussed in section VII.3, the harmonic divergence is strongly correlated with the plasma density scale length L . Therefore, the optimal defocusing distance that maximizes the angular separation ratio is also strongly dependent on L . For smoother plasma-vacuum interfaces (i.e. larger L), the laser wavefronts need to be strongly curved in order to compensate for the important PM focusing. On the other hand, for steep plasma-vacuum interfaces, the PM focusing effect is weak. Thus, the optimal defocusing distance is also small. Figure VII.13 shows the evolution of the ratio $\frac{\Delta\theta}{\theta_n}$ for the 20th harmonic order for different plasma density scale lengths as a function of the laser defocusing distance. The optimal defocusing length is growing as L gets larger. Moreover, the best separation quality ratio seems to be independent of the plasma density scale length. Indeed, as the PM focusing is fully inhibited by the laser wavefront curvature, the plasma scale length L is irrelevant for the harmonic divergence. This proves that this technique is well suited for Doppler HHG separation as it does not require to trade off the HHG efficiency for the attosecond light pulses separation.

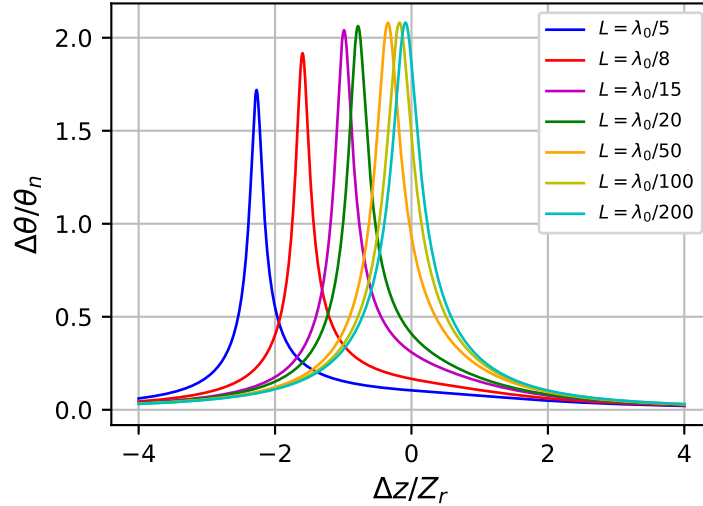


Figure VII.13: Evolution of the ratio $\Delta\theta/\theta_n$ for the 20th harmonic order for three values of scale lengths as a function of the laser defocusing length.

The dependency between the optimal laser defocusing distance and the plasma density scale length is investigated numerically. We performed a 2D PIC parameter scan with respect to the plasma scale length L and the defocusing distance Δz . Given the cost of a single numerically-converged 2D simulation, this numerical study only covers 5 plasma density scale lengths values

$$L \in \left\{ \frac{\lambda_0}{8}, \frac{\lambda_0}{15}, \frac{\lambda_0}{20}, \frac{\lambda_0}{30}, \frac{\lambda_0}{50} \right\}$$

In each case, we performed numerous (between 4 and 8) simulations with different laser defocusing distances, usually near the theoretical optimum predicted by the toy model. In total, we performed 32 2D PIC simulations for a total cost of about half a million core hours on the THETA cluster.

For each simulation, we compute the harmonic angular profile for harmonic orders between 20 and 30. Whenever possible, (when the train is sufficiently separated), the separation ratio $\frac{\Delta\theta}{\theta_n}$ is calculated and the best separation distance that maximizes this ratio is deduced for each plasma density scale length.

Figure VII.14 shows the variation of the angular separation quality with the laser defocusing distance for different plasma density scale lengths. We clearly observe that the optimal defocusing distance increases when the PM density scale length L is increased. We also note that the best numerical separation ratio obtained for different plasma scale lengths does not depend on the value of L and peaks at $\frac{\Delta\theta}{\theta_n} \simeq 2$, in agreement with our theoretical predictions. This suggests that we generate diffraction-free high order harmonics in all cases. The drop in the best separation ratio observed for the case $L = \frac{\lambda_0}{50}$ (red curve) may be explained by a decrease in the harmonic beam size at the PM in this specific case, which increases the harmonic diffraction-free divergence.

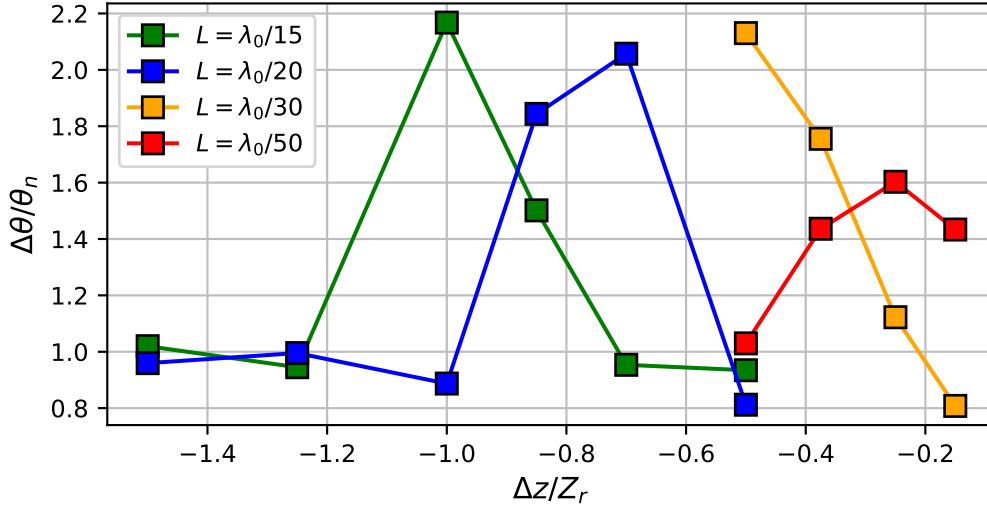


Figure VII.14: Evolution of the separation quality as a function of the defocusing length for different plasma density scale lengths for a 45 degree incidence angle.

A comparison between the best defocusing lengths from the PIC simulations against the ones obtained from the toy model is given in figure VII.15. One can note that our toy model successfully predicts the optimal defocusing distances for every plasma density scale length.

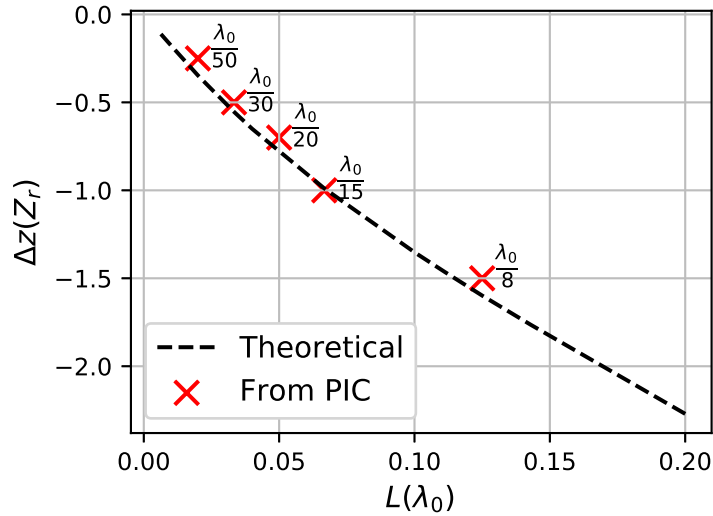


Figure VII.15: Theoretical vs numerical optimal defocusing lengths obtained from 2D PIC simulations.

Figure VII.16 illustrates the angularly resolved harmonic spectra (upper panels), the spatio-temporal profile of the attosecond pulses train far from the PM (or harmonic orders between 20 and 30) (middle panels), and the harmonic angular profiles (lower panels) obtained from PIC simulations involving optimal separation ratios for three plasma scale lengths ($L = \frac{\lambda_0}{8}, L = \frac{\lambda_0}{15}, L = \frac{\lambda_0}{20}$).

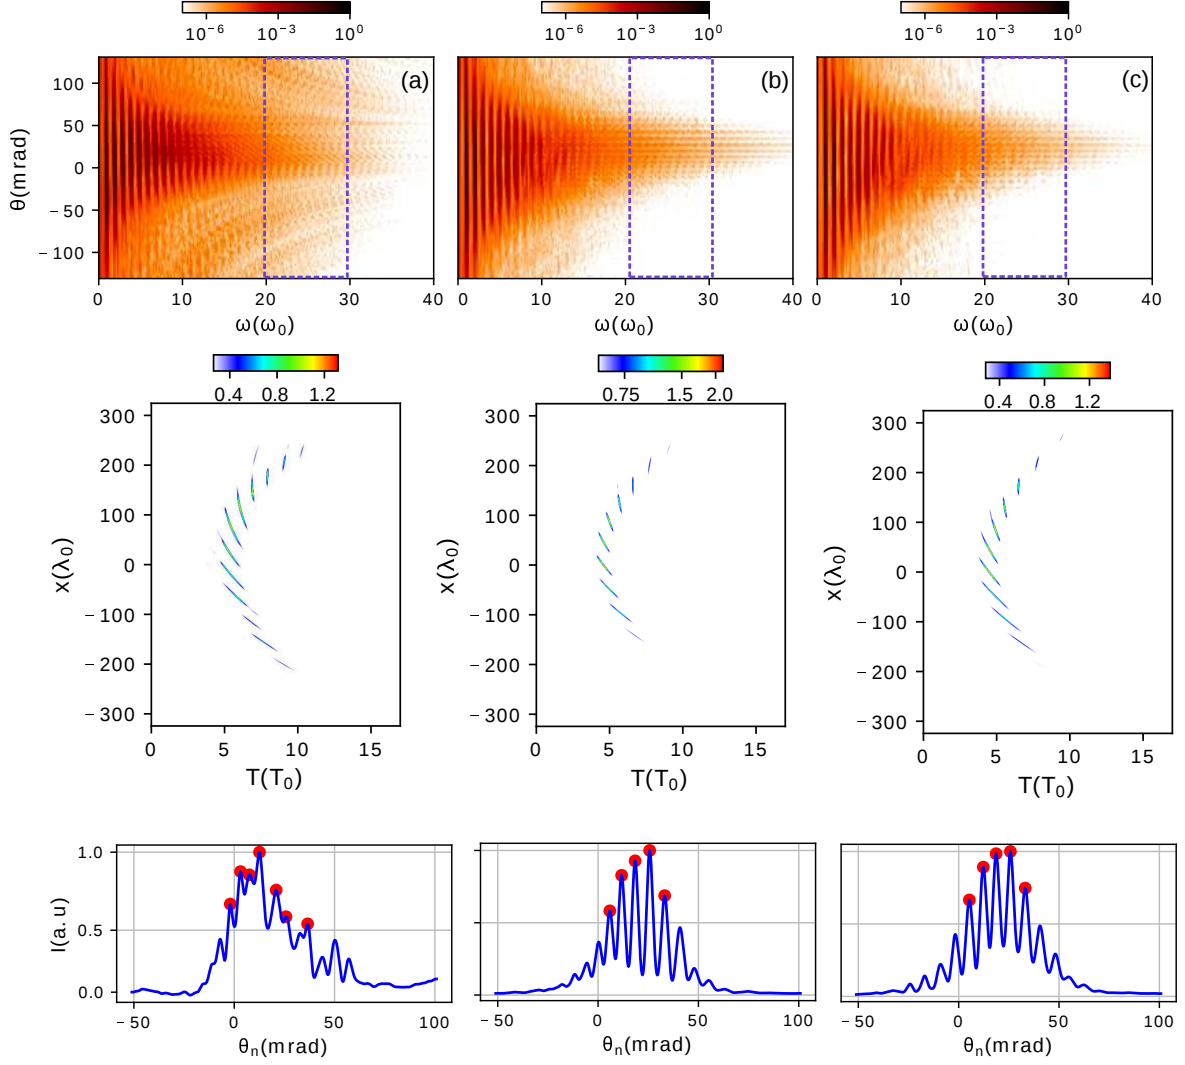


Figure VII.16: Simulation results obtained with the best defocusing distances for different plasma scale lengths. Panel (a): $\lambda_0/8$, $\Delta z = -1.5Z_r$. Panel (b): $\lambda_0/15$, $\Delta z = -Z_r$. Panel (c): $\lambda_0/20$, $\Delta z = -0.7Z_r$. Upper row: angularly resolved harmonic spectra. Central row: attosecond light pulses composed of harmonic orders between 20 and 30 propagated far from the target ($D = 100Z_r \simeq 4000\mu m$). Lower row: angular profiles for the same harmonic range.

The angularly resolved harmonic spectra exhibit high HHG cut-offs in all cases, suggesting a strong HHG efficiency. These profiles are spectrally integrated between harmonic orders 20 and 30 to retrieve the harmonic angular profiles (lower panels). The shape of these angular profiles suggests that a good separation is achieved in the far-field. In the middle panels, we represent the spatio-temporal profiles of attosecond pulses train (for the same harmonic range) far from the PM. We clearly observe that the pulses composing the attosecond train are effectively separated in the far-field, showing that our harmonic divergence reduction technique enables the production of isolated attosecond pulses.

VII.4.5 Effect of the laser incidence angle

Another means to reduce the Doppler harmonic divergence (aside from reducing the plasma scale length) is by increasing the laser incidence angle on target. For instance, by doing so, the laser-plasma interaction region increases in size. Thus the plasma surface is dented on a larger transverse portion. At the same time, the normal component of the laser radiation pressure responsible for the PM denting

drops. Those two effects are schematically illustrated in figure VII.17. This results in a decrease of the maximal PM denting δ_p and an increase in the dented region transverse size (parallel to the plasma surface). Hence, the PM focal distance $f_p \cos(\theta)$ is increased for larger incidence angles and scales as $\propto \frac{1}{\cos(\theta)}$.

The increase in the PM focal distance mitigates the focusing effect on high order harmonics as the maximum quadratic spatial phase ϕ_p (see equation III.63) induced by the PM curvature is reduced. Hence, to compensate the focusing effect of the PM at larger incidence angles, a smaller defocusing distance Δz is required in this case.




Figure VII.17: Schematic representation of the effect of the laser angle of incidence on the PM curvature. The red ellipses represent the driving laser beam. The yellow arrows sketch the radiation pressure normal to the target that induces the PM denting. The arrow sizes schematically represent the radiation pressure magnitude. For oblique incidence (panel (b)), the laser radiation pressure is reduced due to the increase in the size of the interaction region, and the decrease in the normal component of the electromagnetic force exerted on the plasma surface, resulting in a larger PM focal distance (the green dot), and a flatter PM curvature.

To assess this phenomenon, we performed a set of 2D PIC simulations at 55° incidence angle with different laser defocusing distances for two plasma density scale lengths ($L = \lambda_0/15$ and $L = \lambda_0/20$). We want to compare the results obtained in this case with the ones obtained for $\theta_0 = 45^\circ$. All other numerical and physical parameters are the same as those used in the previous simulations (see table VII.1).

In figure VII.18, the red curves represent the evolution of the separation ratio $\frac{\Delta\theta}{\theta_n}$ as a function of the defocusing distance Δz for $L = \lambda_0/15$ (panel (a)) and $L = \lambda_0/20$ panel (b). The blue curves in each panel sketch the evolution of the separation ratio for a 45 degree angle of incidence (taken from figure VII.14).

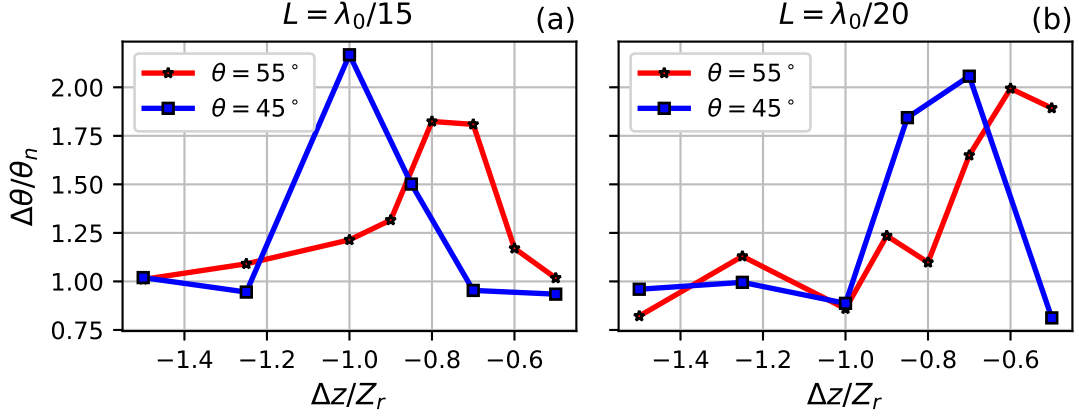


Figure VII.18: Angular separation ratio as a function of the laser defocusing distance: panel (a): $L = \lambda_0/15$, panel (b): $L = \lambda_0/20$. Red curves $\theta = 55^\circ$. Blue curves $\theta = 45^\circ$.

We observe the same dependency pattern between the optimal defocusing distances and the plasma density scale length for both angles of incidence. Besides, by comparing the results obtained for a 55° angle of incidence and the ones for 45° , one can clearly see that the optimal spatial phase compensation is achieved for shorter defocusing distances when the laser incidence angle is larger, which agrees with our analysis of the PM denting. However, this finding does not necessarily result in an improvement of the separation ratio of attosecond light pulses as for both incidence angles, we produce diffraction-free attosecond pulses with minimal divergence.

VII.4.6 Effect of the pulse-front tilt on the angular separation of attosecond light pulses

Up to now, we fixed the laser pulse-front tilt value ξ so that the laser wavefront rotation velocity is maximized at focus. However, counter-intuitively, we now show that this is not the PFT maximizing angular separation.

In panel (a) of figure VII.19, we sketched the evolution of the laser wavefront rotation velocity as a function of the laser defocusing length for 3 values of pulse-front tilt values: ($\xi = 0.5\xi_{max}$ in red, $\xi = \xi_{max}$ in blue, and $\xi = 2\xi_{max}$ in magenta). The laser wavefront rotation velocity is maximum at focus (for any value of pulse-front tilt). However, as the PFT is increased beyond ξ_{max} , the drop rate of WFR as a function of the defocusing distance is slower. This behavior is explained by the increase in the effective laser Rayleigh length Z_r^{eff} for larger pulse-front tilts (due to the increase in the effective laser waist w_{eff}).

Hence, for a fixed non-null defocusing distance Δz , the maximum laser wavefront rotation velocity is achieved by a pulse-front tilt ξ which is always superior to ξ_{max} . This is shown in panel (b) where we sketched the laser wavefront rotation velocity as a function of the pulse-front tilt at different distances from focus ($\Delta z = 0$ in blue, $\Delta z = -Z_r$ in red and $\Delta z = -2Z_r$ in magenta). In each case, the wavefront rotation velocity reaches its peak for a different pulse-front tilt value ξ . This value increases as the defocusing distance is increased.

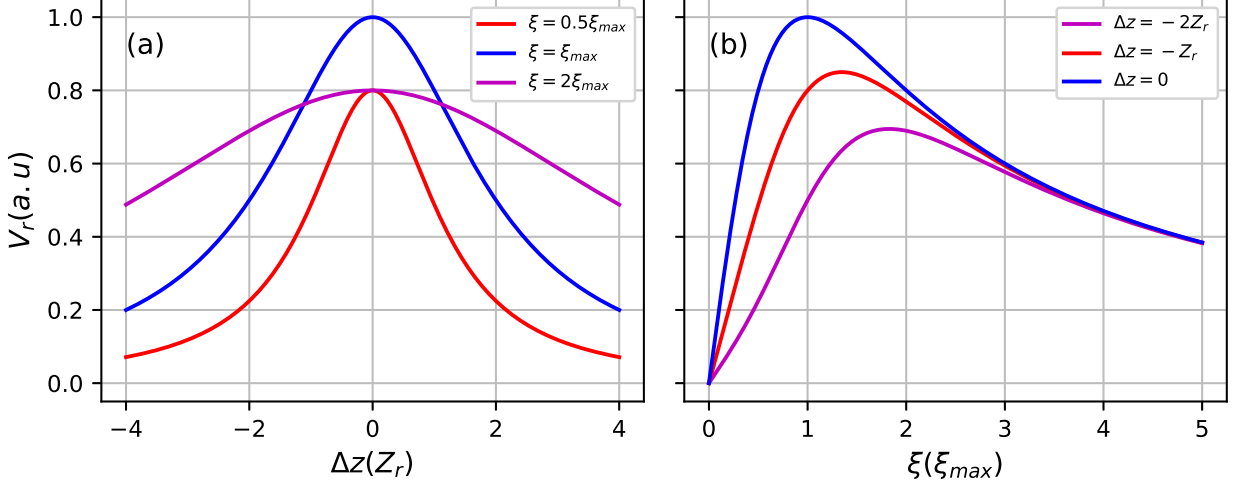


Figure VII.19: Panel (a): Evolution of the laser wavefront rotation as a function of the laser defocusing for multiple values of pulse-front tilts ξ . Panel (b): The wavefront rotation velocity as function of ξ at different locations.

Therefore, choosing pulse-front tilt amounts beyond the critical value $\xi = \xi_{max}$ should result in an increased wavefront rotation velocity at the plasma surface, which is beneficial for the angular separation of attosecond light pulses.

On the other hand, tuning the PFT amount results in an increase of the effective laser waist. Assuming that the harmonic beam sizes scale with the effective laser waist, the resulting harmonics from this setup should be less divergent. Nevertheless, this effect can be mitigated by the drop in the laser intensity on target as the divergence decreases when the laser intensity decreases.

We can use our toy model to assess the effect of the pulse-front tilt parameter on the separation of attosecond light beams. However, our toy model cannot predict the harmonic beam sizes when the laser intensity on target drops. Nevertheless, a qualitative analysis based on these analytical calculations could bring important insights to fetch optimal parameters for attosecond pulses separation.

In figure VII.20, panel (a) sketched the evolution of the angular separation ratio $\frac{\Delta\theta}{\theta_n}$ computed by our analytical model as a function of the laser defocusing length and the pulse-front tilt (for 55° , $L = \frac{\lambda_0}{15}$ and $n=20$). We can note that the best achievable separation ratio is reached for very high laser pulse-front tilt amounts ($\xi \simeq 4.3\xi_{max}$ in this case). We think that this behavior is not physical since our toy model does not take into consideration the effect of the decrease in the laser intensity at the surface of the PM as it is implicitly assumed that the relative harmonic beam sizes at the PM surface is constant ($\frac{w_n}{w_0(\xi, \Delta z)} = 0.8$)

In panel (b), we sketch the evolution of $\frac{1}{\theta_n}$ for the 20th harmonic order (in arbitrary units). The dashed red line outlines each defocusing distance Δz with the pulse-front tilt amount ξ that minimizes the harmonic divergence θ_n through the spatial phase compensation process. We note that the optimal

defocusing distance increases (in absolute value) with respect to the pulse front tilt amount ξ . Indeed, the laser beam size increase induced by the spatial chirp results in an increase in the effective laser Rayleigh length. Therefore, to attain the same wavefront curvature, the laser needs to be defocused over a larger distance and the laser wavefront and PM curvatures matching occurs further from the laser focus when the pulse-front tilt is increased. Moreover, when the spatial phase compensation is achieved (along the green dashed line), the harmonic divergence is reduced for larger spatial chirps amounts ξ due to the decrease in the harmonic diffraction-free divergence.

In panel (c), we represent the evolution of the laser wavefront rotation velocity V_r as a function of ξ and Δz in arbitrary units. The green dashed line outlines each defocusing distance Δz with the pulse-front tilt amount ξ that maximizes the laser wavefront rotation velocity. As discussed earlier, the optimal pulse-front tilt amount increases with respect to the laser defocusing distance. However, for each pulse-front tilt value, the laser wavefront rotation velocity peaks at focus ($\Delta z = 0$).




Figure VII.20: Panel (a): Evolution the angular separation ratio $\frac{\Delta\theta}{\theta_n}$ as a function of the laser defocusing length and the pulse-front tilt ξ . Panel (b): the evolution of $1/\theta_n$ (in arbitrary units) as a function of the laser defocusing length Δz and the pulse-front tilt ξ . The dashed red line highlights the pulse-front tilt that minimizes the harmonic divergence at a given defocusing distance Δz . Panel (c): the evolution of the wavefront rotation velocity as a function of the pulse-front tilt amount and the defocusing distance. The dashed green line highlights the pulse-front tilt amount that maximizes the laser wavefront rotation velocity for a given defocusing distance Δz . In all panels, we assumed a laser incidence angle of 55° and a plasma density scale length of $L = \lambda_0/15$.

In figure VII.21, we extracted the dashed red and green lines from panels (b) and (c) from the previous figure VII.20. The intersection of the two curves takes place at $(\xi_{int}, \Delta z_{int}) = (1.42\xi_{max}, -1.14Z_r)$. For this value of the laser defocusing distance and pulse-front tilt, the PM focusing effect is fully compensated, and at the same time, the laser wavefront rotation velocity reaches it peak with respect to ξ . It might be very interesting to assess the separation efficiency of the lighthouse mechanism in the setup given by $(\xi_{int}, \Delta z_{int})$.




Figure VII.21: Extracted dashed red and green curves from figure VII.20.

Direct observation of the effect of the pulse-front tilt amount on the separation of attosecond light pulses

To study the effect of the pulse-front tilt on the attosecond pulses separation, we performed the following numerical experiment. We fixed the laser defocusing distance to $\Delta z = -Z_r$. We ran a series of 2D PIC simulations for multiple values of ξ between $0.8\xi_{max}$ and $2\xi_{max}$ and two angle of incidence $\theta = 45^\circ$ and $\theta = 55^\circ$. Due to the increase in the laser beam size because of the pulse-front tilt, the simulation box size should be scaled as $\sim \sqrt{1 + (\frac{\xi}{\xi_{max}})^2}$ in 2D. Consequently, these simulations can be extremely expensive in terms of computational time. So we slightly reduced the spatial-temporal resolution for these numerical runs in order to mitigate their total computational cost.

The numerical/physical setup used in this parametric scan is summarized in the following table VII.4:

simulation	Laser params						Plasma params		Num params	
	a_0	θ	FWHM	w_0	$\xi(\xi_{max})$	Δz	L	n_0	dx	ppc
set 1	30.	45°	16fs	$3.2\mu m$	0.8..2.	$-Z_r$	$\lambda_0/15$	$220n_c$	$\lambda_0/190$	6
set 2	30.	55°	16fs	$3.2\mu m$	0.8..2.	$-Z_r$	$\lambda_0/15$	$220n_c$	$\lambda_0/190$	6

 Table VII.4: Numerical/physical parameters for the parametric scan ξ vs θ_0 for a constant defocusing distance.

Figure VII.22 sketches the harmonic angular profiles for three values of pulse-front tilt for each incidence angle between harmonics orders 15 and 20. The best angular separations are achieved with pulse-front tilts values of $\xi = \xi_{max}$ and $\xi = 1.5 \times \xi_{max}$ for $\theta = 45^\circ$ and $\theta = 55^\circ$ respectively. In the first case, the laser defocusing distance $\Delta z = -Z_r$ is optimal for a pulse-front tilt of $\xi = \xi_{max}$. By increasing ξ , the optimal defocusing distance is also increased, and the plasma mirror focusing effect is no longer completely hindered. Despite the increase in the WFR velocity, the harmonic divergence is increased, and attosecond light pulses are no longer separated. On the other hand, the same defocusing distance of $\Delta z = -Z_r$ is above the optimal defocusing distance for 55° (which is $-0.8 \times Z_r$ according to the analytical model). In this case, by increasing the pulse-front tilt, the laser wavefront rotation

velocity is increased and the defocusing distance used in the simulation gets closer to the optimal defocusing distance for attosecond pulses separation, improving the harmonic angular separation.




Figure VII.22: Integrated angular spectrum as a function of ξ and the incidence angle θ for a fixed defocusing distance of $\Delta z = -Z_r$.

Improving the angular separation by tuning the spatial chirp amount

We want to assess if the strategy consisting in simultaneously tuning the laser wavefront curvature and the pulse-front tilt could improve the separation quality of attosecond light pulses.

We performed a 2D PIC simulation with a setup close to that given by figure VII.21. Namely, in this particular case, the $(\xi, \Delta z)$ parameters simultaneously compensate for the PM curvature and maximize the laser wavefront rotation velocity where the laser-plasma interaction takes place. The physical and numerical parameters for this simulation are summarized in the next table VII.5:

Physical parameters								Num params	
Laser parameters					Plasm params			dx	ppcell
a_0	θ	FWHM	w_0	$\xi(\xi_{max})$	$\Delta z(Z_r)$	L	n_0	$\lambda_0/285$	6
30	55°	$16fs$	$3.2\mu m$	1.45	-1.15	$\lambda_0/15$	$220n_c$	$\lambda_0/285$	6

Table VII.5: Numerical/physical parameters simultaneously achieving optimal defocusing distance Δz and laser wavefront rotation velocity.

Figure VII.23 sketches the results obtained from this simulation. In panel (a) we represent the angularly resolved harmonic spectrum of the reflected field. The finger-like structures are clearly visible over the 15th harmonic order. In panel (b), we integrated this spectrum over the harmonic orders H(20-30). The separation ratio is $\frac{\Delta\theta}{\theta_n} = 2.26$ which is slightly better than that obtained for 55° and $\xi = \xi_{max}$ ($\frac{\Delta\theta}{\theta_n} = 1.9$), which proves that tuning the pulse-front tilt along with the defocusing distance

might be beneficial to increase the angular separation of attosecond light pulses.

Finally, in panel (c), we represented the attosecond pulses associated to these harmonic orders at a distance of $D = 4000\mu\text{m} \simeq 100Z_r$ to observe the attosecond pulses separation in the far-field.

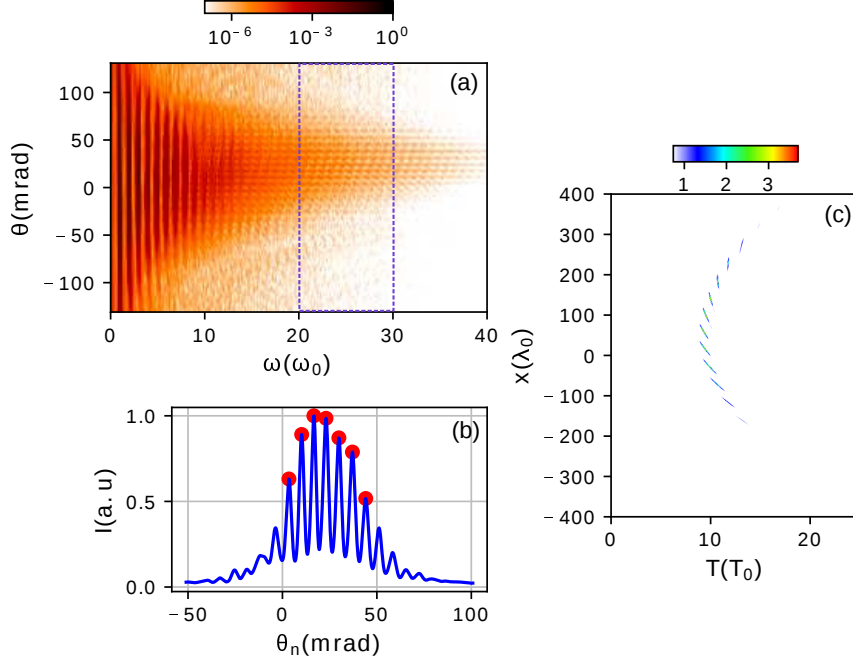


Figure VII.23: Simulation results for an incidence angle $\theta_0 = 55^\circ$, $\xi = 1.45\xi_{max}$ and $\Delta z = -1.15Z_r$.

Attosecond pulses separation with a multi-terawatt laser

In the previous numerical studies, we showed that the attosecond lighthouse effect can be triggered with PetaWatt lasers ($a_0 = 30$). With such a powerful laser, the intensity of the driving laser field at the plasma surface remains extremely high despite the intensity depletion due to the pulse-front tilt/laser defocusing. Hence, in the previous simulations, the Doppler harmonic cutoff, as well as the harmonic beam sizes at the plasma surface, were relatively high, which allows to select higher, and naturally less divergent harmonics. In the ~ 100 TeraWatt regime, the laser intensity is still sufficiently high to drive a relativistic high order harmonic generation process, but with lower harmonic cut-offs and roll-offs.

In this section, we show that our technique also leads to an acceptable separation ratio. To do this, we performed a simulation with the same parameters as for VII.4.3 except that $L = \frac{\lambda_0}{15}$, $\Delta z = -Z_r$ and a normalized laser amplitude $a_0 = 10$. The laser intensity reaches $7 \times 10^{19} \text{W.cm}^{-2}$ at the plasma surface. This setup can be achieved by a few hundreds TeraWatt laser.

The results of this simulation are presented in figure VII.24. Panel (a) illustrates the angularly resolved harmonic spectrum of the reflected field. As one can note, the HHG has significantly dropped compared to the same simulation performed with a higher laser intensity (see figure VII.16 pane (b)). In panel (b), we represent the harmonic angular profile for harmonic orders between 20 and 26. In this configuration, and due to the decrease in the laser intensity, the relativistic gyromagnetic effect is significantly inhibited: multiple attosecond light pulses are emitted in the angular domain $\theta < 0$,

corresponding to the rear of the driving laser field. Furthermore, the most intense attosecond pulse is emitted in the specular direction of the driving laser field ($\theta = 0$ in the streaking probe). By computing the separation ratio obtained by isolating the three central attosecond pulses, we find $\frac{\Delta\theta}{\theta_n} = 1.3 > 1$, which shows that the separation criterion is still fulfilled, although it has decreased compared to the case $a_0 = 30$ where a separation ratio $\frac{\Delta\theta}{\theta_n} = 2$ was achieved. Note also that the central attosecond pulse (emitted along the specular direction) is about two times more intense than its neighboring attosecond pulses. The selected harmonic range is relatively close to the HHG cutoff (cf panel (a)). Hence, the HHG dependency on the instantaneous laser intensity is highly non-linear, and harmonic generation is driven more efficiently by the central laser optical cycle. This relatively high contrast should, in principle, improve the temporal contrast of the isolated attosecond light pulse in the far-field (if the specular attosecond pulse is selected).

The spatio-temporal of attosecond pulses train pertaining to harmonic order between 20 and 26 far from the PM is sketched in panel (c) to illustrate the angular separation of attosecond pulses in the far-field.

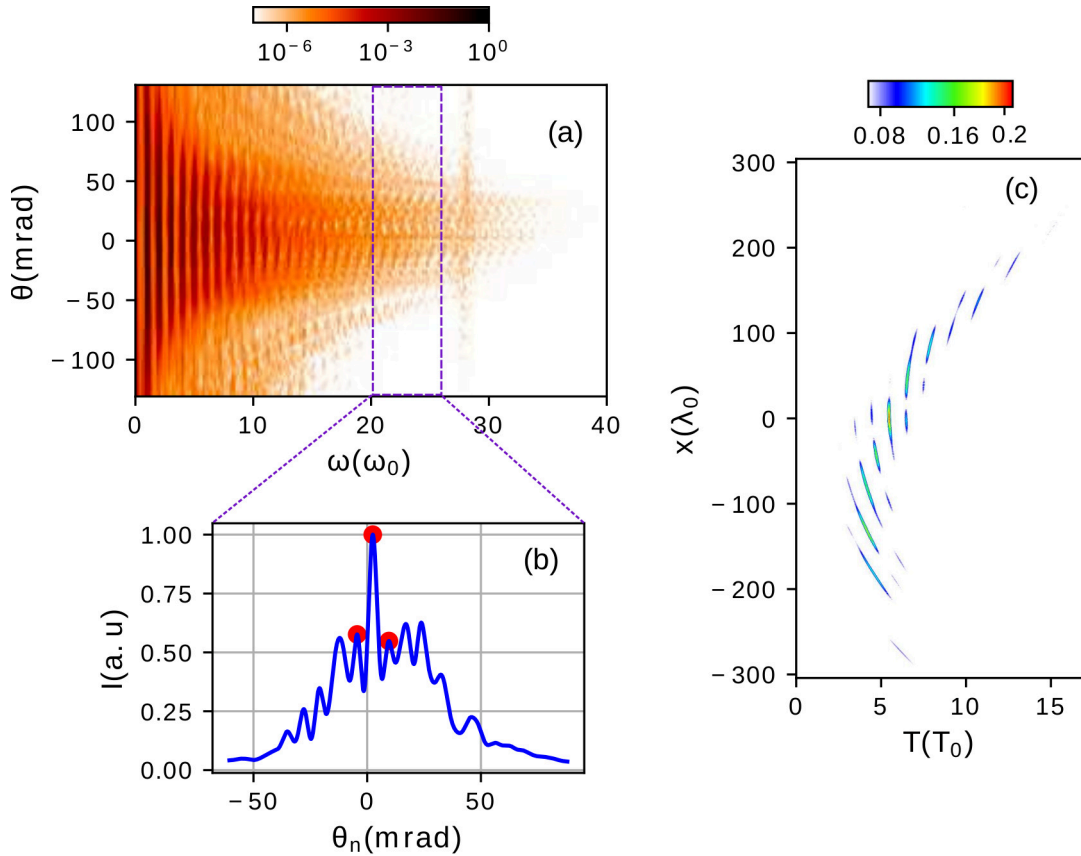


Figure VII.24: Panel (a): Angularly resolved harmonic spectrum of the reflected field. Panel(b): Integrated angular spectrum in the harmonic range between 20 and 26. Panel(c): Propagated attosecond light beams formed by these harmonics over a distance of $D = 4000\mu\text{m} \simeq 100Z_r$.

VII.4.7 Attosecond light pulses separation in 3D

So far, we heavily relied on 2D PIC simulations to investigate the effect of the laser wavefront curvature on the harmonic divergence in the presence of wavefront rotation. In the light of these

results, we were able to validate the accuracy of our harmonic divergence model for a broad range of laser and plasma parameters (plasma scale length, incidence angle, pulse-front tilt, laser intensity). The next step consists in studying the angular separation of attosecond pulses in a full 3D geometry. This study intends to completely validate the efficiency of our harmonic divergence reduction with a 3D PIC simulation and investigate some aspects of the HHG emission that are not accessible with 2D simulations. Such aspects include harmonic astigmatism and harmonic cross section ellipticity.

Physical/numerical setup

For this goal, we performed a single 3D PIC simulation of the lighthouse effect using WARP+PICSAR framework. The physical and numerical parameters are summarized in the next table VII.6:

Physical parameters						Num params			
Laser parameters					Plasm param		dx	ppcell	
a_0	θ	FWHM	w_0	$\xi(\xi_{max})$	$\Delta z(Z_r)$	L	n_0		
40	45°	$22fs$	$2.4\mu m$	1	-1	$\lambda_0/15$	$220n_c$	$\lambda_0/190$	1

Table VII.6: Physical/numerical parameters for the 3D-PIC lighthouse simulation.

The laser duration is longer than for the previously presented results. We checked with the theoretical model and 2D PIC simulations that for a 22 fs laser, a good angular separation is achievable. Additionally, the laser beam size has been reduced to $2.4\mu m$ in order to shrink the simulation box and save computation time. One particle per cell was used for both electrons and ions. Note that the numerical electric charge of each pseudo particle varies between $\pm 130\bar{e}$ and $\pm 28000\bar{e}$ (with \bar{e} the elementary electron charge) when the plasma density varies between n_c and $n_0 = 220n_c$.

Figure VII.25 is a schematic representation of the simulation box in 3D:




Figure VII.25: Schematic representation of the simulation box for the 3D simulation. The laser antenna (red rectangle) is used to input a laser beam with WFR. The streaking probe (purple rectangle) captures the reflected field after the laser-plasma interaction. Note that the laser transverse profile is elliptical, with a larger diameter along the chirped axis.

The PM surface is orthogonal to the x -axis. The laser is p-polarized (the transverse magnetic field is along y -axis) and its wavefronts are rotating in the x - z plane (around the y axis). The reflected field is captured by a streaking probe located at a position of $\simeq 10\lambda_0$ from the target surface.

This simulation was run on 32,768 BLUE GENE-Q nodes of the MIRA supercomputer at ALCF for 20 hours and cost more than 10 millions core hours.

We perform the same post-processing operations on the reflected field as in 2D, namely:

1. Back propagate the reflected field from the streaking plane to the PM surface.
2. Compute the 3D angularly resolved harmonic spectrum of the harmonic field $B_y(\theta_x, \theta_y, \omega)$.
3. Compute the harmonic angular profile $I(\theta_x, \theta_y, \omega_{min}, \omega_{max})$.

Parallel post-processing implementation

Performing the 3D post-processings is very computationally demanding on its own. The reflected field occupies hundreds of gigabytes of memory, which can easily exceed the total memory of a single compute node. Therefore, to be able to perform the data analysis of this massive simulation, we developed specific parallel post-processing tools to parallelize all the computations involved in the data analysis.

The field propagation and filtering require FFT computations and point-to-point multiplication of 3D arrays. If the data are distributed in memory, these two operations can be parallelized using the adequate tools: namely a distributed-memory FFT library. We use mpi4py-FFT [105] library to perform distributed memory FFT computations with Python. This library supports both the pencil and the slab decompositions detailed in V.3.1. The block multiplication in Fourier space can be performed

serially on each compute node and involves no data communications.

The computation of the angularly resolved harmonic spectrum $I(\theta_x, \theta_y, \omega)$ is achieved by performing a variable substitution in the harmonic field Fourier spectrum $|B_y(k_x, k_y, \omega)|^2$. This variable substitution reads:

$$\begin{aligned}\theta_x &= \arcsin\left(c\frac{k_x}{\omega}\right) \Leftrightarrow k_x = \sin(\theta_x)\frac{\omega}{c} \\ \theta_y &= \arcsin\left(c\frac{k_y}{\omega}\right) \Leftrightarrow k_y = \sin(\theta_y)\frac{\omega}{c}\end{aligned}\tag{VII.16}$$

Performing this computation in parallel requires a specific data layout that is represented in figure VII.26. A slab decomposition is used to decompose the Fourier spectrum of the harmonic beam along the frequency axis ω . With this setup, each processor is appended to a harmonic subdomain $\omega_{min}, \omega_{max}$, while the k_x and k_y are locally stored in each processor's memory. This way, the variable substitution given by equation VII.16, which computationally translates into an interpolation in the discretized Fourier space, is a local operation and does not require any communication between processing units.

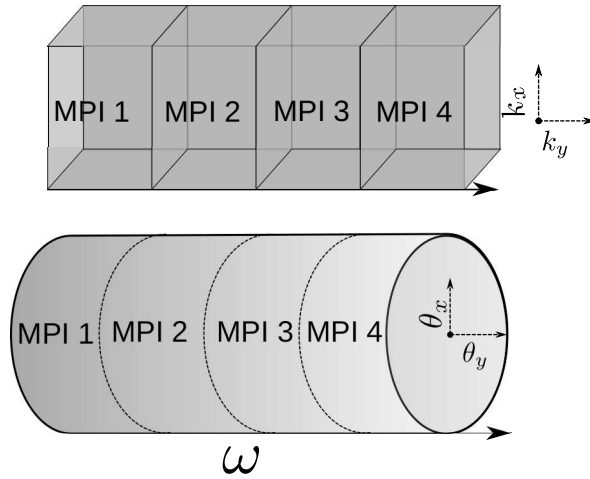


Figure VII.26: Data layout used for parallelizing the computation of the angularly-resolved spectrum.

Angularly resolved harmonic spectrum

Figure VII.27 represents the far-field angular profile for 4 different harmonic ranges, specified in the title of each panel.

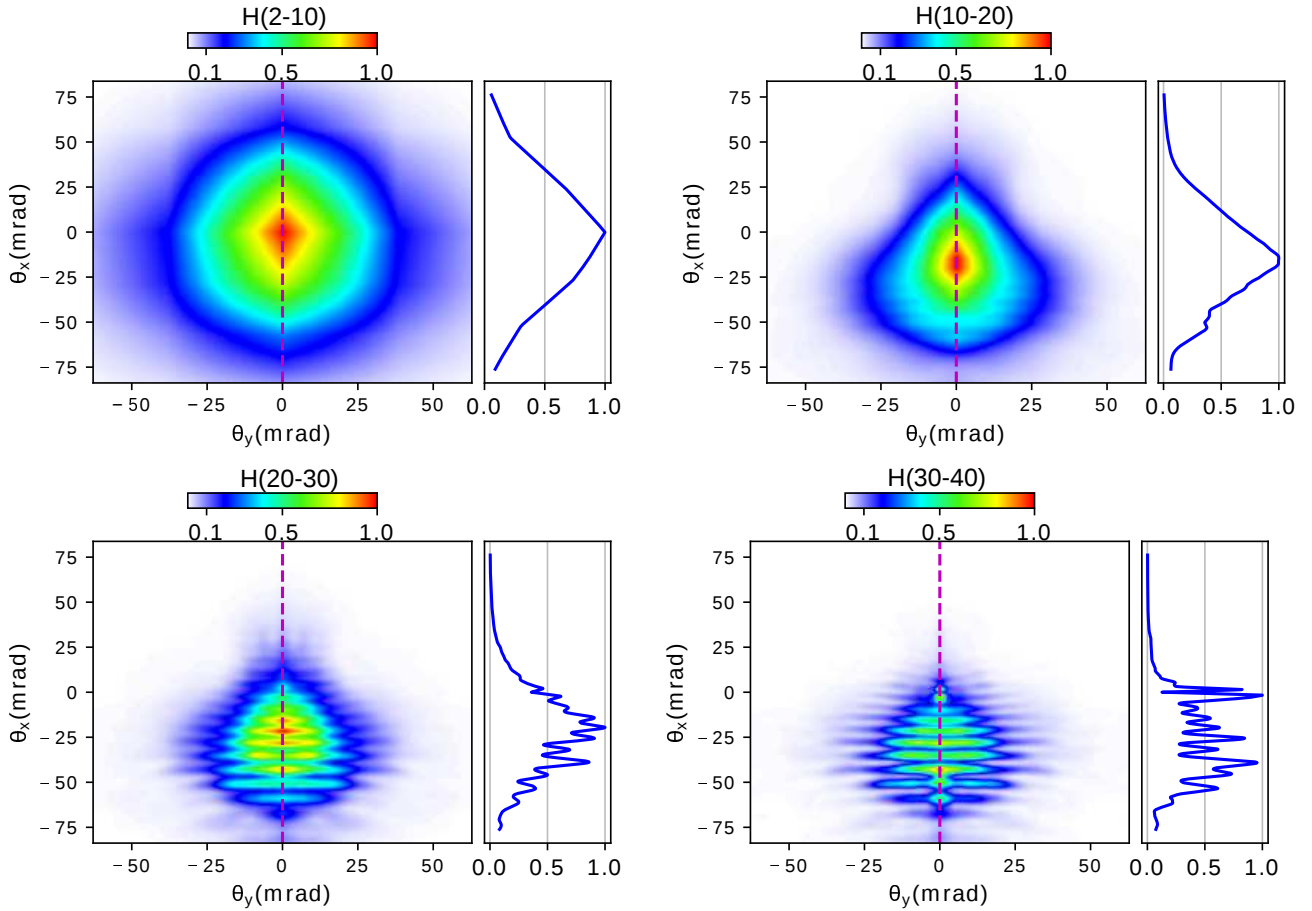


Figure VII.27: Angularly resolved harmonic spectrum integrated over 10 harmonic orders. Each panel represents a different harmonic range. Each spectrum is normalized to its maximum value. The blue curves next to each panel display the evolution of the angular spread along the chirped axis along which the attosecond pulses separation is occurring, for $\theta_y = 0$.

Discussion

First, one can note that attosecond light pulses are angularly separated for a very high harmonic range (above the 20th harmonic order). This was expected as the laser duration in this simulation was longer than the one used in the 2D parameter scans, thus the wavefront rotation velocity is slower.

Moreover, we can observe from panels (c) and (d) that the harmonic divergence is larger for the y-axis than the x-axis. This is emphasized in figure VII.28 where we extracted the angular profile of a unique attosecond light pulse from panel (d). This means that the attosecond light pulses are elliptically shaped in the far-field, with a longer diameter (higher divergence) along the non chirped y-axis. This was expected as the laser cross section is elliptically shaped at the PM surface, with a long diameter along the chirped axis x, thus the cross section of attosecond light pulses is also elliptical with the same ellipticity as the laser (cf figure VII.29 panel (c)). The ellipticity of these pulses will invert as they propagate far from the PM surface, and the shortest diameter (along y-axis) at the PM surface will induce the largest divergence in far-field.

Nevertheless, the laser cross section ellipticity at the PM surface could only induce a maximum harmonic divergence ellipticity ratio between x and y axes of $\sqrt{2}$. By performing a least square fit

with Gaussian functions on the data of figure VII.28, we find

$$\frac{\theta_y^{num}}{\theta_x^{num}} = \frac{17.6}{4.2} = 4.2 \gg \sqrt{2}$$

This can be explained by the fact that, the defocusing distance used in this simulation $\Delta z = -Z_r$ induces a laser wavefront curvature that matches the PM curvature in the x-z plane solely (the plane where the laser wavefronts are rotating). Nevertheless, the laser curvature induced along the non chirped y-axis over-matches the PM curvature in the y-z plane. Indeed, as shown by the simulation specified by table VII.2 (cf figure VII.10), the optimal defocusing distance that exactly compensates for the PM focusing effect in the absence of wavefront rotation and for a plasma density scale length of $L = \frac{\lambda_0}{15}$ and for an incidence angle of 45° is $\Delta z(\xi = 0, L = \lambda_0/15, \theta_0 = 0) = -0.6Z_r > -Z_r$. Therefore, not only the PM focusing effect is fully inhibited along y axis, but also, they are emitted with a non-null defocusing quadratic spatial phase. As a result, the high order harmonics produced by this divergence reduction technique are astigmatic: while the attosecond pulses focal spot along the chirped axis exactly coincides with the PM surface thanks to the spatial phase compensation technique, the focal spot along y-axis is located somewhere in the PM bulk (cf panels (a) and (b) of figure VII.29). We can use the harmonics divergence model (i.e equation VII.15) to compute the divergence for the same harmonic range along the two transverse axes respectively. To compute θ_y , we supposed an oblique incidence at 45° and no spatial chirp ($\xi = 0$). We obtain a divergence ratio

$$\frac{\theta_y^{model}}{\theta_x^{model}} = 5.5$$

, which is in good agreement with the measured ratio of 4.2, which shows that our harmonic divergence model is robust in 3D geometry.




Figure VII.28: A zoomed image over one attosecond light pulse from panel(d) of figure VII.27. The angular spread along y axis is larger than the angular spread along x axis.

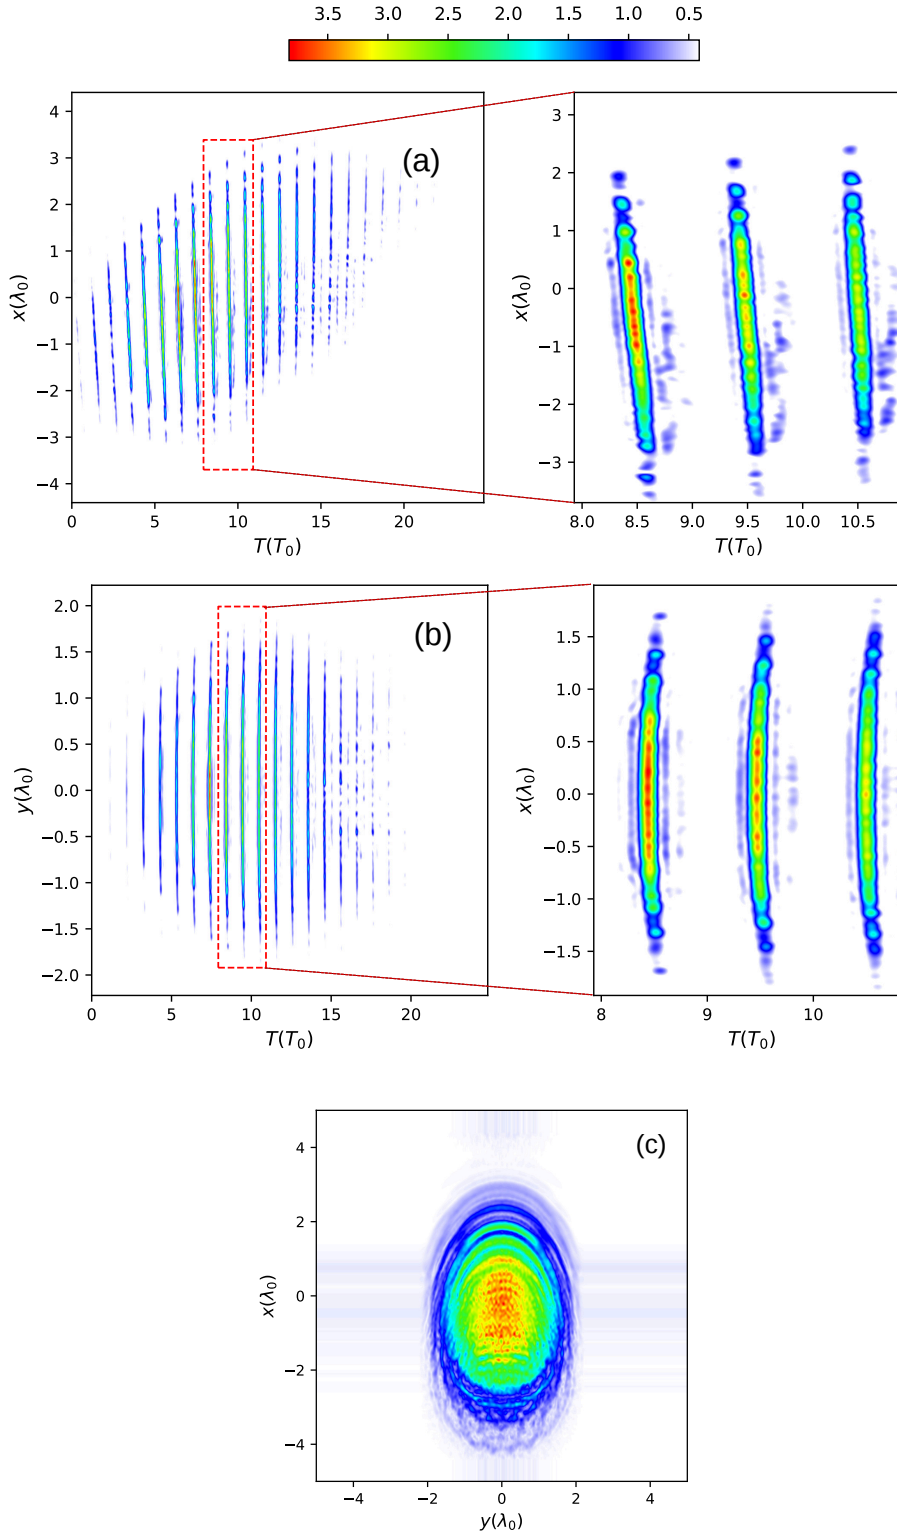


Figure VII.29: Different slices of the attosecond pulses train for harmonic orders between 20 and 30 at the PM surface. Panel (a): Slice in the x - z plane, in the middle of the y -axis. The side panel zooms over three successive attosecond pulses. As one can note, the wavefront of these pulses are flat, which shows that the spatial phase compensation has been achieved. Panel (b): Slice in the y - z axis, in the middle of the x -axis. The side panel zooms over the same three attosecond pulses as in Panel (a). This time, the attosecond wavefronts are curved, this curvature corresponds to an over-compensated PM curvature. Panel (c): Transverse profile of an attosecond pulse at the PM surface, exhibiting an elliptical shape.

Finally one can note that there is little to no attosecond light beams emitted in for $\theta_x > 0$ due to

the strong gyromagnetic effect induced by these ultra relativistic laser intensities. In figure VII.30, we represented the evolution of the filtered laser field, and the attosecond pulses train for H-20-30 at the central point of the field streaking. As one can notice, the harmonic emission efficiency significantly drops before reaching the laser central optical cycle, suggesting that the plasma surface dynamics is significantly distorted to a point where the PM no longer radiates high order harmonics.




Figure VII.30: The blue curve represents the evolution of time envelopes of the attosecond light beams composed of H(20-30), emitted at center of the field streaking. The red curve is the filtered reflected laser field, also from the center of the streaking.

VII.5 Laser pulse shaping

In this section, we propose another technique that can be combined with the lighthouse mechanism to produce isolated attosecond light pulses. Just like the spatial phase compensation technique presented in the previous section, this new approach aims at mitigating the divergence of individual attosecond light pulses produced by a spatially chirped laser beam irradiating a plasma mirror surface.

VII.5.1 General principle

The technique presented in this section is based on a controlled distortion of the laser beam transverse profile in order to appropriately manipulate the PM curvature and inhibit the focusing effect it exerts on the emitted high order harmonics. The general principle of this technique is schematically illustrated in figure VII.31. In the relativistic regime of interaction, the laser radiation pressure dents the PM surface within the laser-plasma interaction region. The shape of the dented plasma depends on the laser profile. For a Gaussian laser beam, the laser intensity variations induce a parabolic PM curvature, represented in panel (a), which in return strongly focuses the high order harmonics (cf panel (c)). On the other hand, if the laser transverse profile was flat (cf panel (b)), the radiation pressure should homogeneously dent the PM surface. Therefore, the high order harmonics generated within such a setup should be less divergent as they are less subject to the focusing effect induced by the PM curvature (panel (d)). Coupled with the lighthouse mechanism, this technique could lead to the emission of unfocused, and angularly separated attosecond light pulses.

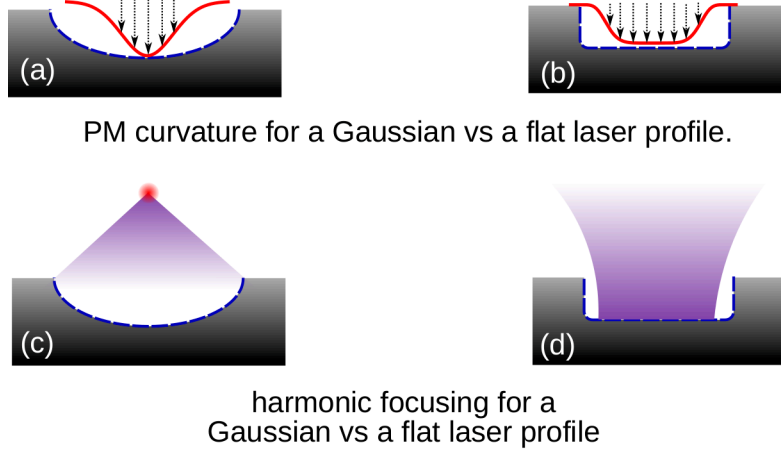


Figure VII.31: Panel (a): Typical PM curvature under the effect of a Gaussian laser beam. Panel (b): Plasma mirror denting for a flat or an hyper Gaussian laser transverse profile. Panel (c): harmonic focusing by a curved plasma surface. Panel (d): harmonic propagation with a flat PM surface.

To investigate the effectiveness of this idea, we need to spatially manipulate the transverse profile of the driving laser beam. In the following subsection, we show how to leverage the laser spatial chirp to control its transverse profile.

VII.5.2 Flattening a laser pulse profile in the presence of spatial chirp

In the presence of spatial chirp, the laser central frequency ω_0 varies as a function of the transverse position at focus x_f . This dependency is illustrated in figure VII.32. Panels (a) and (c) represent an electromagnetic field at focus in the presence of wavefront rotation in the (x_f, ω) and (x_f, t) spaces respectively. The laser frequency dependency on the transverse position is emphasized in panels (b) and (d) where we represented three transverse slices of the electromagnetic fields from panels (a) and (c) respectively: as one can note, due to the wavefront rotation, the temporal gap between two successive laser wavefronts is squeezed for $x_f > 0$ and expanded for $x_f < 0$ (pane d), resulting in a shorter apparent laser wavelength for $x_f < 0$ than for $x_f > 0$.

This particular aspect of the wavefront rotation coupling suggests that the laser cross section could be shaped by acting on the laser Fourier spectrum $E(:, \omega)$. In practice, such distortion can be performed by using a Dazzler device (see references [106, 107] for more information). This device is usually used to modify the laser temporal profile by acting on its Fourier spectrum. In the presence of spatio-temporal couplings, it could also be used to shape the laser spatial profile.

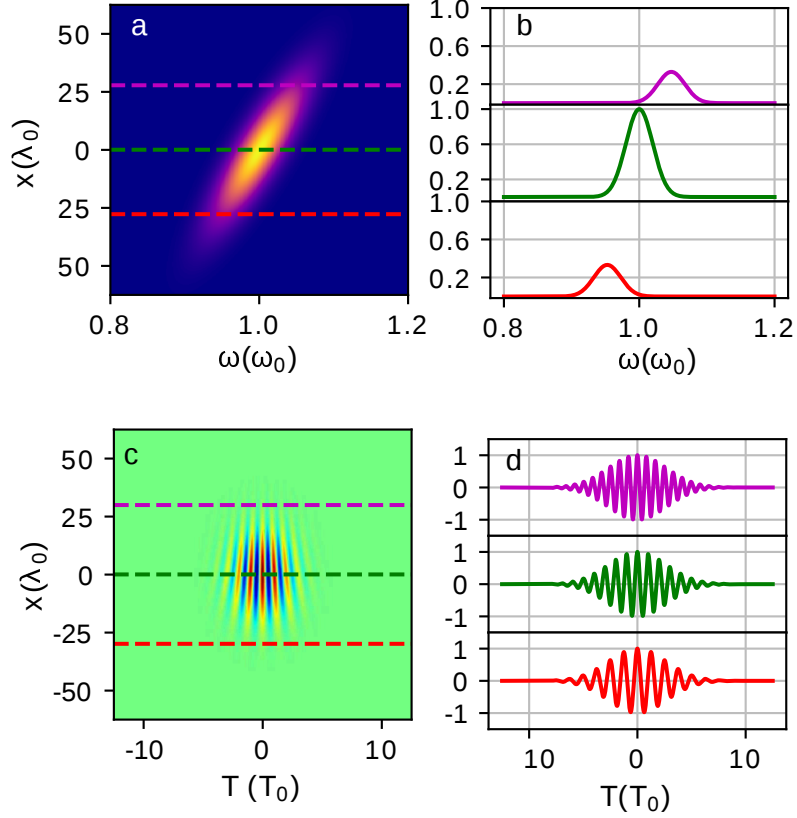


Figure VII.32: Panel(a): Spatio-spectrally chirped electromagnetic field in the spatio-spectral space (x, ω) . Panel(b): 3 slices of the field at different transverse positions. Panel(c): The same field in the spatio-temporal space (x, t) with (panel(b)) three slices of the field at different transverse positions.

Concretely, homogenizing the laser beam cross section requires to damp the laser intensity at the center of the focal spot ($|x_f/w_0| \ll 1$). Such a distortion could be achieved by damping the laser Fourier spectrum around the central laser frequency ω_0 . Due to the spatial-chirp, the laser intensity is solely mitigated at the center of the focal spot, where the central laser frequency is focused, which should eventually result in a more homogeneously distributed laser intensity across the focal spot.

Numerically, the damping process is emulated by applying a filter on the laser beam in the Fourier space. In our simulations, we used a filter with a gain function $G(\omega)$ given by:

$$G(\omega) = 1 - c_1 \exp \left[- \left((\omega - \omega_0) \times \frac{\tau_0}{c_2} \right)^2 \right] \quad (\text{VII.17})$$

c_1 and c_2 are tuning parameters that are used to control the amplitude as well as the standard deviation of the filter gain function. After filtering, the laser energy loss is compensated to conserve the total laser energy.

We tried different values of c_1 and c_2 with the aim of flattening the laser pulse shape. For a pulse-front tilt $\xi = \xi_{max}$ we noted that there is no efficient set of parameters (c_1, c_2) that effectively homogenizes the laser intensity. Indeed, for $\xi = \xi_{max}$, the damping filter affects large portion of the laser cross section profile, resulting in an inefficient shaping. We found out that an efficient laser cross

section flattening is possible for larger pulse-front tilt amounts ($\xi > \xi_{max}$). The reason behind this is related to the spectral spread over the focal spot associated to the spatial chirp. In the (x_f, ω) space, the electromagnetic beam at focus is given by :

$$\hat{E}(x_f, \omega) \propto \exp\left[-\left(\frac{x_f}{w_{eff}}\right)^2\right] \times \exp\left[-\left(\frac{\tau_{eff}}{2}\right)^2(\omega - \omega_c(x_f, \xi))\right] \quad (\text{VII.18})$$

where $\omega_c(x_f, \xi)$ is the laser central frequency as a function of the transverse position x_f and is given by:

$$\omega_c(x_f, \xi) = \omega_0 + \frac{2}{\tau_0} \frac{\xi}{\xi_{max}} \frac{x_f}{w_{eff}} \quad (\text{VII.19})$$

We define the total spectral spread over the focal spot as the difference between the central frequencies from both sides of the focal spot $x_f \pm w_{eff}$.

$$\begin{aligned} \Delta\Omega(\xi) &= |\omega_c(w_{eff}, \xi) - \omega_c(-w_{eff}, \xi)| \\ \Delta\Omega(\xi) &= \frac{4}{\tau_0} \left| \frac{\xi}{\xi_{max}} \right| \end{aligned} \quad (\text{VII.20})$$

This quantity grows as a function of the pulse-front tilt. Therefore, by increasing ξ , the spectral content of the laser focal spot (between $-w_{eff}$ and w_{eff}) is stretched (i.e. more frequencies are focused within this spatial extent). Consequently, the damping filter $G(\omega)$ given by equation VII.17, which acts on the frequency range $\omega_0 \pm \frac{\tau_0}{c_2} \omega_0$ would distort the laser transverse profile on a smaller fraction of the total cross section size, leading to a more efficient shaping.

Figure VII.33 summarizes the effect of this filtering on the laser spatio-spectral and spatio-temporal profiles. For this test, we used a spatial chirp associated to a pulse-front tilt of $\xi = 1.5\xi_{max}$ and a waist of $w_0 = 3.2\mu m$. The filter tuning parameters are $c_1 = 0.7$, $c_2 = 1.3$. These parameters will be used later on for simulations. In panel (a), we represent the laser profile in the (x_f, ω) space after the spectral shaping. Because the laser profile is tilted, the filtering will only affect transverse positions close to $x_f = 0$. In panel (b), we sketched three slices of the field profile in the (x_f, ω) from panel (a) for 3 different transverse positions : $x_f = 0.75w_{eff}$ (blue curve), $x_f = 0$ (green curve), $x_f = 0.75w_{eff}$ (red curve). The black dashed line represents the filter gain function as a function of the frequency ω . The central slice (in green) is significantly damped by the filter to a point where the laser amplitude for the three transverse position is roughly the same. Panel (c) illustrates the spatio-temporal shape of the filtered laser beam in (x, t) space. As one can note, the field spatio-temporal profile is significantly disrupted at the rear and the front of the beam. However, at the center of the beam, the cross section is well flatted, as one can see in panel (d) where we represented the laser intensity profile for $t = 0$ in blue. In the red dashed line, we sketched the laser intensity profile without the pulse shaping process. This pulse shaping process leads to a decrease in the maximum laser intensity at focus by about 30% despite the total pulse energy conservation. This drop is however acceptable as the laser cross section has been effectively flattened in the end. The parameters used in this test will be employed to conduct 2D PIC simulations in order to assess the effectiveness of this technique in producing low divergence and separated attosecond light pulses.

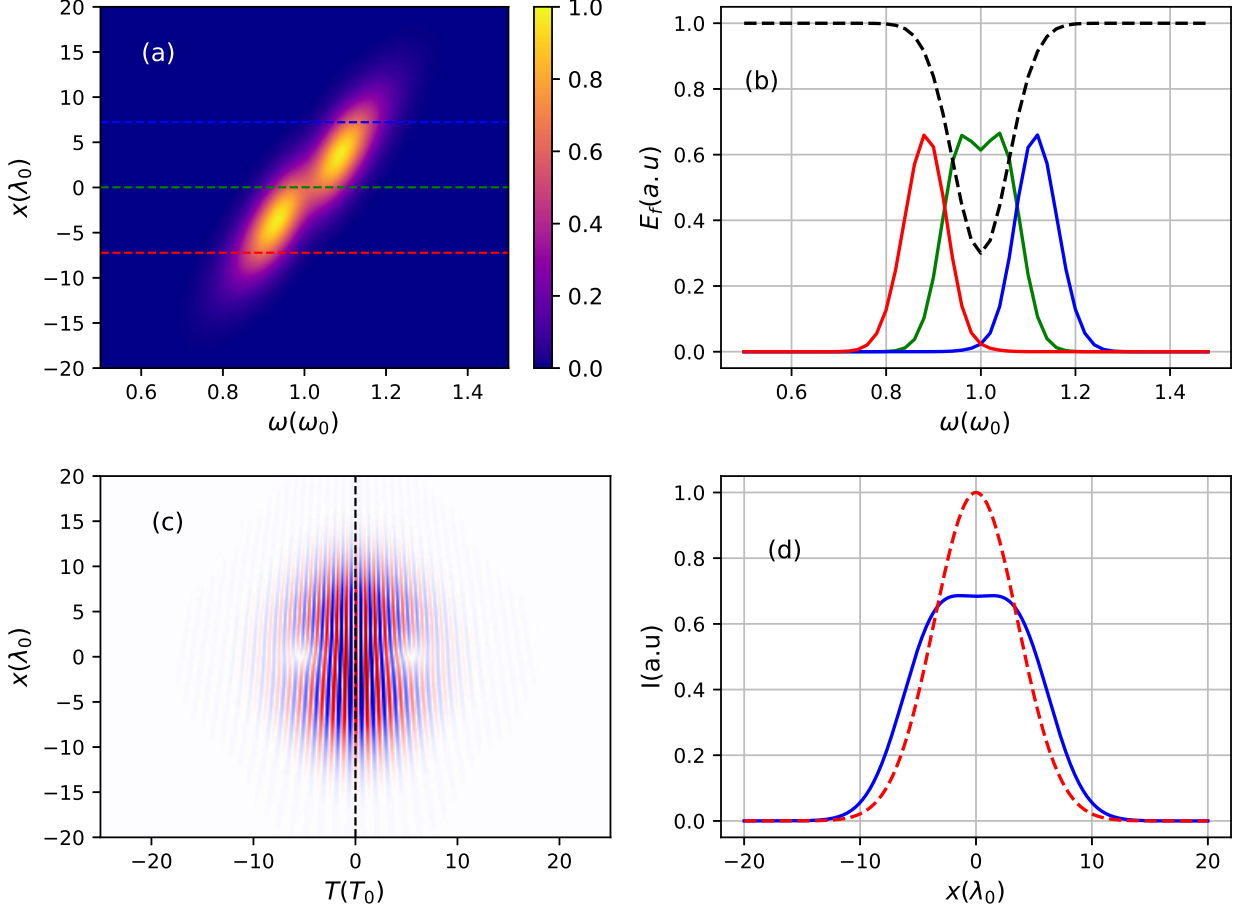


Figure VII.33: Panel (a): Laser profile in spatio-spectral space (x, ω) after filtering. Panel (b): Three longitudinal slices of the field in the spatio spectral space for $x_f = 0$ in green, $x_f = 0.75w_{eff}$ in blue and $x_f = -0.75w_{eff}$ in red. Panel (c): spatio-temporal shape of the shaped laser beam. Panel (d): the blue curve represents the laser intensity profile at the center of the laser pulse. The red curve is the intensity profile obtained in the absence of the laser pulse shaping.

VII.5.3 PIC simulations of ROM harmonics generation with a shaped laser pulse

The laser pulse shaping process has been implemented in the WARP+PICSAR PIC code. The laser is injected in the simulation by an electric current antenna that induces the desired laser spatio-temporal profile. Given the laser parameters, (τ, w_0, ξ, a_0) , the whole laser beam profile is computed in the (x, t) space. Then, it is Fourier transformed to the (k, ω) space and multiplied by the filter gain function $G(\omega)$ given by equation VII.17. Before going back to the spatio-temporal space, the field is retro-propagated by a distance equal to the distance between the plasma surface and the antenna location via a plane-wave decomposition and the electromagnetic energy loss is compensated. This way, the shaped laser field is perfectly focused at the plasma target and it transports the same amount of energy as a non shaped pulse. At each time step of the PIC loop, the laser antenna injects a slice of the filtered field obtained by the procedure described here.

The effect of the laser cross section shaping on the separation of attosecond light pulses has been investigated using 2D PIC simulations. Unfortunately, this technique has not been tested in a full

3D geometry given the expensive computational cost of such simulations. Nevertheless, 2D PIC simulations are sufficient to test the new approach and compare it with the laser defocusing technique presented previously.

For this matter, we performed three 2D PIC simulations whose parameters are summarized in the following table VII.7:

sim	Laser parameters								Plasm params		Num params	
	a_0	θ	FWHM	w_0	$\xi(\xi_{max})$	c_1	c_2	$\frac{\Delta z}{Z_r}$	L	n_0	dx	ppc
sim 1	30	55°	16fs	$3.2\mu m$	1.5	None	None	0	$\frac{\lambda_0}{20}$	$220n_c$	$\frac{\lambda_0}{285}$	6
sim 2	30	55°	16fs	$3.2\mu m$	1.5	None	None	-1	$\frac{\lambda_0}{20}$	$220n_c$	$\frac{\lambda_0}{285}$	6
sim 3	30	55°	16fs	$3.2\mu m$	1.5	0.7	1.3	0	$\frac{\lambda_0}{20}$	$220n_c$	$\frac{\lambda_0}{285}$	6

Table VII.7: Physical and numerical setups for simulations used to compare the laser pulse shaping technique with the defocusing technique and the standard lighthouse scheme.

- *sim 1* is the original lighthouse scheme. We performed this simulation to compare the original basic separation quality of the lighthouse scheme with our newly developed divergence reduction techniques.
- *sim 2* exploits the laser wavefront curvature to compensate for the PM focusing effect. We will use the data from this simulation to compare both separation techniques.
- Finally, *sim 3* is the setup where the laser cross section is shaped and focused at the plasma surface.

Figure VII.34 sketches the far-field angular profiles obtained from each of these simulations and pertaining to the harmonic range (15-20). As expected, without any additional separation optimization, the attosecond lighthouse effect by itself is not sufficient to produce angularly separated light pulses in the relativistic regime (panel (a)). Panel (b) illustrates the angular separation achievable by the laser defocusing technique. Panel (c) sketches the results obtained with the new pulse shaping scheme. As one can note, a very good separation is also achieved in this case. By comparing *sim 2* and *sim 3*, we can notice that for $\theta_x > 50$ mrad, no attosecond pulses separation is occurring, while multiple pulses are already radiated separately in *sim2*. This is probably due to the distortion in the laser forefront induced by our shaping. By comparing the attosecond pulses separation from *sim2* and *sim3*, one can note that the new method produces more angularly spaced but more divergent attosecond pulses. The increase in the angular shift between successive pulses was anticipated: by focusing the laser pulse on the PM surface, the laser wavefront rotation velocity is higher than that for a defocused laser (such as in *sim 2*). The increase in the attosecond pulses divergence shows that the focusing effect of the PM is not completely mitigated by the laser transverse shaping approach compared to the laser defocusing technique.

All in all, by computing the separation ratio obtained from the simulations *sim 2* and *sim 3*, we find $\frac{\Delta\theta}{\theta_n} = 1.72$ for *sim 2* and $\frac{\Delta\theta}{\theta_n} \simeq 1.89$ for *sim 3*. Hence, a slightly better separation is achieved with this new technique.




Figure VII.34: Integrated angular spectrum for each simulation case from table VII.5.3.

In figure VII.35, we represent the attosecond pulses train associated to harmonic orders between 15 and 20 obtained from *sim 3*. We can see that during the first laser optical cycles, no separation is occurring. However, multiple attosecond light pulses clearly separate few laser optical cycles later on.

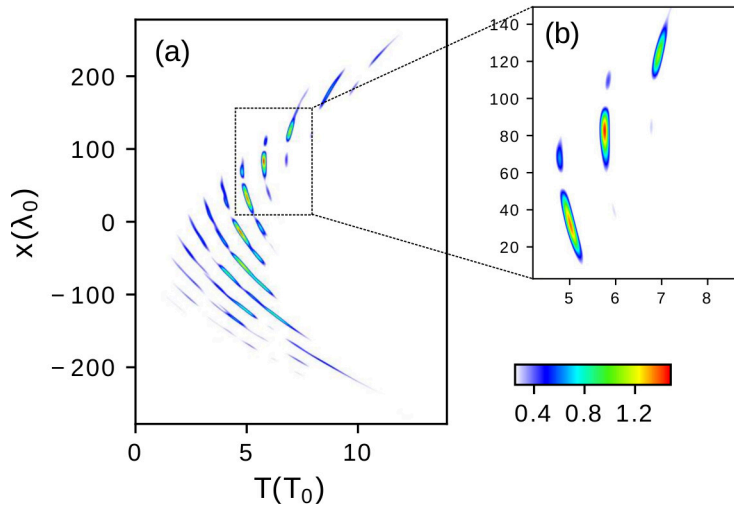


Figure VII.35: Attosecond pulses obtained from *sim 3* propagated for $100 Z_r$ away from the plasma surface.

VII.6 Conclusion

In this chapter, we have shown that it is possible to enable an efficient separation of ROM attosecond light by combining the attosecond lighthouse effect with harmonic divergence reduction techniques.

We have identified two approaches to mitigate the divergence of attosecond light pulses in the relativistic regime of interaction that both result in the splitting of the attosecond pulses train into multiple isolated ultra-fast light bunches. The first technique exploits the laser wavefront curvature to inhibit the focusing effect of the curved PM. We have shown that this scheme can be improved by tuning the pulse-front tilt amount to increase the wavefront rotation velocity out of focus.

The second technique is probably more challenging to implement experimentally as it requires a good control over the laser beam cross section shape. Nevertheless, we were able to develop an efficient strategy to spatially shape the laser pulse by taking advantage of the spectral aspects of the spatio-temporal coupling responsible for the laser wavefront rotation mechanism.

By enabling a good separation of attosecond light pulses, the lighthouse mechanism provides a natural angular mapping of the temporal dynamics of the PM. Hence, it can be used as an in-situ measurement to probe the laser-plasma interaction process in the relativistic regime temporally. In this context, by analyzing the divergence of single attosecond light pulses as well as the angular spacing between successive light spikes constituting the attosecond pulses train, one could simultaneously retrieve the PM denting velocity as well as the femtosecond chirp of the attosecond pulses train. Finally, producing isolated attosecond light pulses in the ROM regime may pave the way to unprecedented capabilities in the field of ultra-fast spectroscopy and attosecond sciences, by providing a source of extremely bright X-UV attosecond pulses.

Chapter VIII

Conclusions and perspectives

The first part of this work is dedicated to the conception and development of a novel massively parallel strategy for the pseudo-spectral FFT-based Maxwell's equations solvers. This type of solvers is necessary for a broad range of PIC-UHI simulations that require a dispersion-less electromagnetic field description. However, such approaches usually come with limited scaling performances and important memory footprints at very large scales. The innovative parallelization strategy exposed in this manuscript leverages the multi-level parallelism of modern super-computers to achieve unmatched performances in terms of both strong scaling and memory use. This technique exploits the moderate scalability of distributed memory FFT algorithms along with standard domain decomposition techniques to minimize the total computation time of FFT-based Maxwell solvers while significantly reducing the total memory occupancy of fields arrays. Coupled with the WARP+PICSAR PIC code, this achievement is essential to tackle large scales 3D PIC simulations exploiting pseudo-spectral Maxwell solvers.

Furthermore, I have implemented a unique parallel diagnostic tool for UHI physics. This tool can be used to reconstruct simulation data in any Lorentz boosted frame on the fly. It has been benchmarked on high order harmonic generation simulations in different regimes. With this tool, it is possible to overcome the complexity underlying the oblique laser incidence in interpreting and analyzing simulations results. In particular, these benchmarks show that the coupling between the plasma dynamics and the high order harmonics properties is better described with this Lorentz boosted frame diagnostic. This tool can also be used in the context of laser-plasma acceleration to reconstruct laboratory frame data when performing the simulation in a Lorentz boosted frame.

This computational work enabled us to conduct an extensive numerical and theoretical study to investigate the feasibility of generating isolated attosecond light pulses from relativistic plasma mirror harmonics. In particular, we were interested in transposing the attosecond lighthouse effect to Doppler harmonics. In this context, attosecond light pulses generated in this regime of interaction are too divergent to be angularly separated with the lighthouse effect due to the strong plasma mirror curvature resulting from the radiation pressure exerted by the driving laser pulse. For this matter, we propose two novel schemes to produce spatially separated attosecond light pulses from plasma

mirrors in the relativistic regime.

Both techniques consist in tailoring the laser spatial phase or amplitude profile to reduce the natural divergence of Doppler harmonic beams. The first one consists in placing the PM out of laser focus so that the laser wavefront curvature compensates for the spatial phase induced by the curvature of the PM. The second one consists in tailoring the laser amplitude profile in order to homogenize the laser intensity distribution across the laser focus. By doing so, the plasma mirror is homogeneously dented across the laser-plasma interaction region, which substantially efficiently inhibits the harmonics wavefronts curvatures.

By identifying optimal conditions, we show that it is possible to produce divergent-free high order harmonics that result in separated attosecond light pulses in the far-field with both techniques. For the first approach, we developed a semi-analytical toy model to predict the angular separation quality of attosecond pulses train. This model was fully corroborated with PIC simulation results in 2D and 3D geometries under various interaction conditions.

The second technique can be achieved by employing a Dazzler device, usually employed to temporally shape the driving laser pulse. In the presence of spatio-temporal couplings, this operation also impacts the laser transverse profile. We identified optimal shaping parameters to efficiently mitigate the plasma mirror curvature, while maintaining optimal harmonic generation conditions. This approach was also corroborated by several 2D PIC simulations.

Perspectives

The novel parallelization strategy presented in this manuscript should render massively parallel 3D PIC simulations for UHI physics considerably more affordable in terms of computational cost. This approach could also be transposed to other numerical solvers used to address Cauchy-type problems involving physical propagations of information at finite speed (e.g. heat equation, Schrödinger equation, Navier-Stokes equations...).

Furthermore, the Doppler attosecond pulses separation schemes proposed in this thesis can be used in plasma-HHG experiments to generate isolated X-UV pulses, which may be useful for ultra-fast spectroscopy experiments or as an in-situ probe for laser-plasma interaction in the relativistic regime. While many important challenges still need to be addressed before enabling the use of high order harmonics from solid targets as a source of ultra-short probing pulses for external applications (i.e. increase the laser stability and repetition rate ...), the spatial mapping of the temporal dynamics of the plasma, enabled by the attosecond lighthouse scheme brings important insights about the physics at play in the context of laser-plasma interaction (e.g. gyrokinetic signature on the harmonic profile, plasma mirror curvature over time ...).

Appendix A

Derivation of laser parameters out of focus in the presence of wavefront rotation

In this appendix, we derive the laser parameters out of focus, in the presence of wavefront rotation **at focus**. We relied on the symbolic computing software Mathematica [108] to carry out the calculations involved in this appendix. In the focal plane, the transverse electric field $E(x, t, \Delta z = 0)$ writes:

$$E(x, t, \Delta z = 0) = E^0 \exp \left[-\frac{x^2}{w_{eff}^2} \right] \exp \left[-\frac{t^2}{\tau_{eff}^2} \right] \exp [i \times t (\omega_0 + \zeta x)] \quad (\text{A.1})$$

with

$$\begin{aligned} \tau_{eff} &= \tau \sqrt{1 + \left(\frac{\xi}{\xi_{\max}} \right)^2} \\ w_{eff} &= w_0 \sqrt{1 + \left(\frac{\xi}{\xi_{\max}} \right)^2} \\ \zeta &= \frac{2\xi w_i}{w_0 \tau_{eff}^2} \end{aligned} \quad (\text{A.2})$$

where ξ is the pulse-front tilt amount converted into WFR by a focusing lens of focal length f . w_i and w_0 are the pulse waists before focus and at focus in the absence of pulse-front tilt respectively. We recall that $w_0 = \frac{\lambda_0 f}{\pi w_i}$. Besides, τ is the pulse duration before focus. Finally, $\xi_{\max} = \frac{\tau}{w_i}$ is the pulse-front tilt amount that maximizes the WFR velocity at focus.

In the space-frequency domain (x, ω) , this field reads:

$$\begin{aligned} \widehat{E}(x, \omega, \Delta z = 0) &= TF_t(E(x, t, \Delta z = 0))(\omega) \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(x, t, \Delta z = 0) \exp(-i\omega t) dt \\ &= E^0 \tau \frac{1}{\sqrt{2}} \exp \left(-\left(\frac{x}{w_{eff}} \right)^2 - \frac{\tau_{eff}^2}{4} (\omega_0 + \zeta x - \omega) \right) \end{aligned} \quad (\text{A.3})$$

A second Fourier transform with respect to x gives:

$$\begin{aligned}
\widehat{E}(k_x, \omega, \Delta z = 0) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp(-ik_x x) \widehat{E}(x, \omega, \Delta z = 0) dx \\
&= E^0 \frac{\tau_{eff} w_{eff}}{\sqrt{4 + \tau_{eff}^2 w_{eff}^2 \zeta^2}} \\
&\times \exp\left(\frac{k_x^2 w_{eff}^2 + \omega \tau_{eff}^2 (-2\omega_0 + ik_x w_{eff}^2 \zeta) - ik_x \omega_0 \tau_{eff}^2 w_{eff}^2 \zeta + \omega^2 \tau_{eff}^2 + \omega_0^2 \tau_{eff}^2}{\tau_{eff}^2 w_{eff}^2 \zeta^2 + 4}\right)
\end{aligned} \tag{A.4}$$

The propagation operation in the Fourier space is given by:

$$\widehat{E}(k_x, \omega, \Delta z = z) = \widehat{E}(k_x, \omega, \Delta z = 0) \exp\left[ik_z(\omega, k_x)z\right] \tag{A.5}$$

where $k_z(\omega, k_x)$ is given by:

$$k_z(\omega, k_x) = \sqrt{\frac{\omega^2}{c^2} - k_x^2} \tag{A.6}$$

Under the slowly varying envelope approximation (in time and space), we have $k_x \ll \frac{\omega_0}{c}$ and $\omega \ll \omega_0$. Therefore, $k_z(\omega, k_x)$ reads:

$$k_z(\omega, k_x) \simeq \frac{\omega}{c} - \frac{k_x^2}{2k_0} \tag{A.7}$$

where $k_0 = \frac{\omega_0}{c}$. Therefore, A.6 becomes:

$$\widehat{E}(k_x, \omega, \Delta z = z) = \widehat{E}(k_x, \omega, \Delta z = 0) \times \exp\left[iz\left(\frac{\omega}{c} - \frac{k_x^2}{2k_0}\right)\right] \tag{A.8}$$

The field in the spatio-temporal space is given by the inverse Fourier transform of equation A.8:

$$\begin{aligned}
E(x, t' = t - z/c, \Delta z = z) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \widehat{E}(k_x, \omega, \Delta z = z) dk_x d\omega \\
&\times \exp\left[\frac{i\left(2\omega_0\left(t'\tau_{eff}^2\left(w_{eff}^2(\omega_0 + x\zeta) + 2icz\right) + it'^2 w_{eff}^2 + i\tau_{eff}^2 x^2\right) - ct'^2 z\left(\tau_{eff}^2 w_{eff}^2 \zeta^2 + 4\right)\right)}{4\tau_{eff}^2 c\left(Z_r^{eff} + iz\right)}\right] \\
&\times \exp\left[it'\left(\omega + \frac{\zeta x}{1 + i\frac{z}{Z_r^{eff}}}\right)\right] \times \exp\left[-\frac{t'^2}{\tau_{eff}^2\left(1 + i\frac{z}{Z_r^{eff}}\right)}\right] \times \exp\left[-\frac{k_0 x^2}{2\left(Z_r^{eff} + iz\right)}\right] \\
&\times \exp\left[-\frac{it'^2 z w_{eff}^2 \zeta^2}{4\left(Z_r^{eff} + iz\right)}\right] \times \exp\left[-i\frac{t'^2 z}{Z_r^{eff} \tau_{eff}^2\left(1 + i\frac{z}{Z_r^{eff}}\right)}\right] \\
&\times \exp\left[it'\left(\omega + \frac{\zeta x}{1 + i\frac{z}{Z_r^{eff}}}\right)\right] \times \exp\left[-\frac{k_0 x^2}{2\left(Z_r^{eff} + iz\right)}\right] \times \exp\left[-\frac{t'^2}{\tau_{eff}^2}\right] \\
&\times \exp\left[-\frac{it'^2 z w_{eff}^2 \zeta^2}{4\left(Z_r^{eff} + iz\right)}\right]
\end{aligned} \tag{A.9}$$

Using the relationships $\zeta = \frac{2\xi w_i}{w_0 \tau_{eff}^2}$, $w_{eff}^2 = w_0^2 \left(1 + \left(\frac{\xi}{\xi_{max}}\right)^2\right)$ and $\xi_{max} = \frac{\tau}{w_i}$, we have:

$$\begin{aligned} w_{eff}^2 \zeta^2 &= 4\xi^2 w_i^2 \left(1 + \left(\frac{\xi}{\xi_{max}}\right)^2\right) \times \frac{1}{\tau^4 \left(1 + \left(\frac{\xi}{\xi_{max}}\right)^2\right)^2} \\ &= \frac{4}{\tau_{eff}^2} \left(\frac{\xi}{\xi_{max}}\right)^2 \end{aligned} \quad (\text{A.10})$$

Therefore:

$$\begin{aligned} \exp\left[-\frac{it'^2 z w_{eff}^2 \zeta^2}{4(Z_r^{eff} + iz)}\right] &= \exp\left[-izt'^2 \frac{\xi^2}{\tau_{eff}^2 \xi_{max}^2 (Z_r^{eff} + iz)}\right] \\ &= \exp\left[-iZ_r^{eff} z t'^2 \frac{\xi^2}{\tau_{eff}^2 \xi_{max}^2 ((Z_r^{eff})^2 + z^2)}\right] \times \exp\left[-t'^2 \frac{\xi^2}{\xi_{max}^2 \tau_{eff}^2} \frac{z^2}{(Z_r^{eff})^2 + z^2}\right] \end{aligned} \quad (\text{A.11})$$

Using equation A.11, we can factorize equation A.9 as follows:

$$\begin{aligned} E(x, t' = t - z/c, \Delta z = z) &\propto \exp\left[it' \left(\omega + \underbrace{\frac{\zeta x}{1 + i \frac{z}{Z_r^{eff}}}}_{\text{STC}}\right)\right] \\ &\times \exp\left[\underbrace{\frac{ik_0 x^2}{2(iZ_r^{eff} - z)}}_{\text{Waist+WFC}}\right] \\ &\times \exp\left[\underbrace{-\frac{t'^2}{\tau_{eff}^2} \left(1 + \frac{\xi^2 z^2}{\xi_{max}^2 [(Z_r^{eff})^2 + z^2]}\right)}_{\text{Temporal envelope}}\right] \\ &\times \exp\left[\underbrace{-it'^2 \frac{\xi^2}{\tau_{eff}^2 \xi_{max}^2 \left(\frac{Z_r^{eff}}{z} + \frac{z}{Z_r^{eff}}\right)}}_{\text{Temporal chirp}}\right] \end{aligned} \quad (\text{A.12})$$

Expression A.12 shows that the laser complex parameter (i.e the term under-braced in blue) is equal to that of a regular Gaussian pulse with a waist of w_{eff} . Hence, the laser beam size and radius of curvature out of focus read respectively:

$$w_{eff}(\Delta z = z) = w_{eff} \sqrt{1 + \left(\frac{z}{Z_r^{eff}}\right)^2} \quad (\text{A.13})$$

and

$$R(\Delta z = z) = -z \left[1 + \left(\frac{Z_r^{eff}}{z}\right)^2\right] \quad (\text{A.14})$$

On the other hand, the propagation of the laser pulse out of focus gives rise to a temporal chirp (under-braced in magenta) that reads:

$$\beta(\xi, \Delta z = z) = -\frac{\xi^2}{\tau_{eff}^2 \xi_{max}^2 \left(\frac{Z_r^{eff}}{z} + \frac{z}{Z_r^{eff}} \right)} \quad (\text{A.15})$$

Moreover, the laser temporal duration out of focus (under-braced in violet) reads:

$$\tau_{eff}^2(\Delta z = z) = \tau_{eff}^2 \frac{\xi_{max}^2 \left[1 + \left(\frac{Z_r^{eff}}{z} \right)^2 \right]}{\xi^2 + \xi_{max}^2 \left[1 + \left(\frac{Z_r^{eff}}{z} \right)^2 \right]} \quad (\text{A.16})$$

Finally, the spatio-temporal coupling term (under-braced in red) reads:

$$a_{xt} = \frac{i\zeta}{1 + i \frac{z}{Z_r^{eff}}} \quad (\text{A.17})$$

Equation A.12 also shows that the propagation in vacuum of a laser pulse exhibiting WFR induces a pulse-front tilt coupling ($\text{Re}(a_{xt}) \neq 0$) out of focus:

$$\begin{aligned} \text{PFT} &= \text{Re}(a_{xt}) \\ \text{PFT} &= \frac{\zeta}{\frac{Z_r^{eff}}{z} + \frac{z}{Z_r^{eff}}} \end{aligned} \quad (\text{A.18})$$

Moreover, the WFR parameter out of focus reads:

$$\begin{aligned} \zeta(\Delta z = z) &= \text{Im}(a_{xt}) \\ \zeta(\Delta z = z) &= \frac{\zeta}{1 + \left(\frac{z}{Z_r^{eff}} \right)^2} \end{aligned} \quad (\text{A.19})$$

The wavefront rotation velocity out of focus is given by:

$$V_r(\Delta z = z) = \frac{d \left(\frac{k_x}{k} \right)}{dt'} \quad (\text{A.20})$$

with $k_x(t', \Delta z = z) = \zeta(\Delta z = z)t'$ and $k(t', \xi, \Delta z = z) = \frac{\omega_0 - t'\beta}{c}$, the instantaneous laser wavevector. In this thesis, we consider laser pulses with spatial chirps associated with $\xi \sim \xi_{max}$. Moreover, the laser defocusing distances of interest are not larger than the Rayleigh length $|z| \lesssim Z_r^{eff}$. Under these conditions, the temporal chirp out of focus is of the same order as the inverse square of the laser temporal duration (see equation A.15):

$$\beta \sim \frac{1}{\tau_{eff}^2}$$

Therefore, the laser frequency variation $\Delta\omega$ due to the temporal chirp is of order:

$$\Delta\omega \sim \tau_{eff} \times \beta \sim \frac{1}{\tau_{eff}} \ll \omega_0$$

(under the slowly varying envelope approximation (i.e $\tau_{eff} \gg T_0 = \frac{2\pi}{\omega_0}$))

Consequently, the temporal chirp out of focus does not affect much the instantaneous laser frequency :

$$k(\xi, \Delta z = z, t') \simeq \frac{\omega_0}{c} = k_0$$

Finally, the laser wavefront rotation velocity out of focus reads:

$$\boxed{V_r(\Delta z = z) = \frac{c\zeta(\Delta z = z)}{\omega_0} = \frac{1}{k_0} \frac{\zeta}{1 + \left(\frac{z}{Z_r^{eff}}\right)^2}} \quad (\text{A.21})$$

Appendix B

Relationship between spatio-spectral beam sizes

	(x, t)	(x, ω)	(k, ω)	(k, t)
$a_{xx} =$	a_{xx}	$b_{xx} - \frac{b_{x\omega}^2}{4b_{\omega\omega}}$	$\frac{1}{4c_{kk} - \frac{c_{k\omega}^2}{c_{\omega\omega}}}$	$\frac{1}{4d_{kk}}$
$b_{xx} =$	$a_{xx} - \frac{a_{xt}^2}{4a_{tt}}$	b_{xx}	$\frac{1}{4c_{kk}}$	$\frac{1}{4d_{kk} - \frac{d_{kt}^2}{d_{tt}^2}}$
$c_{kk} =$	$\frac{1}{4a_{xx} - \frac{a_{xt}^2}{a_{tt}^2}}$	$\frac{1}{4b_{xx}}$	c_{kk}	$d_{kk} - \frac{d_{kt}^2}{4d_{tt}}$
$d_{kk} =$	$\frac{1}{4a_{xx}}$	$\frac{1}{4b_{xx} - \frac{b_{x\omega}^2}{b_{\omega\omega}^2}}$	$c_{kk} - \frac{c_{k\omega}^2}{4c_{\omega\omega}}$	d_{kk}

Table B.1: Relationship between spectro-spatial beam sizes different domains.

	(x, t)	(x, ω)	(k, ω)	(k, t)
$a_{tt} =$	a_{tt}	$\frac{1}{4b_{\omega\omega}}$	$\frac{1}{4c_{\omega\omega} - \frac{c_{k\omega}^2}{c_{kk}}}$	$d_{tt} - \frac{d_{kt}^2}{4d_{kk}}$
$b_{\omega\omega} =$	$\frac{1}{4a_{tt}}$	$b_{\omega\omega}$	$c_{\omega\omega} - \frac{c_{kt}^2}{4c_{kk}}$	$\frac{1}{4d_{tt} - \frac{d_{kt}^2}{d_{kk}^2}}$
$c_{\omega\omega} =$	$\frac{1}{4a_{tt} - \frac{a_{xt}^2}{a_{xx}^2}}$	$b_{\omega\omega} - \frac{b_{x\omega}^2}{4b_{xx}}$	$c_{\omega\omega}$	$\frac{1}{4d_{tt}}$
$d_{tt} =$	$a_{tt} - \frac{a_{tx}^2}{4a_{xx}}$	$\frac{1}{4b_{\omega\omega} - \frac{b_{x\omega}^2}{b_{xx}^2}}$	$\frac{1}{4c_{\omega\omega}}$	d_{tt}

Table B.2: Relationship between spectro-temporal beam sizes in different domains.

Bibliography

- [1] Donna Strickland and Gerard Mourou. Compression of amplified chirped optical pulses. *Optics Communications*, 56(3):219 – 221, 1985.
- [2] A. Dubietis, G. Jonušauskas, and A. Piskarskas. Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in bbo crystal. *Optics Communications*, 88(4):437 – 440, 1992.
- [3] Gerard A. Mourou, Toshiki Tajima, and Sergei V. Bulanov. Optics in the relativistic regime. *Rev. Mod. Phys.*, 78:309–371, Apr 2006.
- [4] E. Esarey, C. B. Schroeder, and W. P. Leemans. Physics of laser-driven plasma-based electron accelerators. *Rev. Mod. Phys.*, 81:1229–1285, Aug 2009.
- [5] A. J. Gonsalves, K. Nakamura, J. Daniels, C. Benedetti, C. Pieronek, T. C. H. de Raadt, S. Steinke, J. H. Bin, S. S. Bulanov, J. van Tilborg, C. G. R. Geddes, C. B. Schroeder, Cs. Tóth, E. Esarey, K. Swanson, L. Fan-Chiang, G. Bagdasarov, N. Bobrova, V. Gasilov, G. Korn, P. Sasorov, and W. P. Leemans. Petawatt laser guiding and electron beam acceleration to 8 gev in a laser-heated capillary discharge waveguide. *Phys. Rev. Lett.*, 122:084801, Feb 2019.
- [6] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J.-P. Rousseau, F. Burgy, and V. Malka. A laser–plasma accelerator producing monoenergetic electron beams. *Nature*, 431(7008):541–544, September 2004.
- [7] C. G. R. Geddes, Cs. Toth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans. High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding. *Nature*, 431(7008):538–541, September 2004.
- [8] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick. Monoenergetic beams of relativistic electrons from intense laser–plasma interactions. *Nature*, 431(7008):535–538, September 2004.
- [9] S. Corde, K. Ta Phuoc, G. Lambert, R. Fitour, V. Malka, A. Rouse, A. Beck, and E. Lefebvre. Femtosecond x rays from laser-plasma accelerators. *Rev. Mod. Phys.*, 85:1–48, Jan 2013.

- [10] C Thaury and F Quéré. High-order harmonic and attosecond pulse generation on plasma mirrors: basic mechanisms. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 43(21):213001, oct 2010.
- [11] U. Teubner and P. Gibbon. High-order harmonics from laser-irradiated plasma surfaces. *Rev. Mod. Phys.*, 81:445–479, Apr 2009.
- [12] R. Lichters, J. Meyer-ter Vehn, and A. Pukhov. Short-pulse laser harmonics from oscillating plasma surfaces driven at relativistic intensities. *Physics of Plasmas*, 3(9):3425–3437, 1996.
- [13] S. V. Bulanov, N. M. Naumova, and F. Pegoraro. Interaction of an ultrashort, relativistically strong laser pulse with an overdense plasma. *Physics of Plasmas*, 1(3):745–757, 1994.
- [14] T. Baeva, S. Gordienko, and A. Pukhov. Theory of high-order harmonic generation in relativistic laser interaction with overdense plasma. *Phys. Rev. E*, 74:046404, Oct 2006.
- [15] C. Thaury, F. Quéré, J.-P. Geindre, A. Levy, T. Ceccotti, P. Monot, M. Bougeard, F. Réau, P. d’Oliveira, P. Audebert, R. Marjoribanks, and Ph Martin. Plasma mirrors for ultrahigh-intensity optics. *Nature Physics*, 3:424 EP –, Apr 2007. Article.
- [16] F. Quéré, C. Thaury, P. Monot, S. Dobosz, Ph. Martin, J.-P. Geindre, and P. Audebert. Coherent wake emission of high-order harmonics from overdense plasmas. *Phys. Rev. Lett.*, 96:125004, Mar 2006.
- [17] M. Thévenet, A. Leblanc, S. Kahaly, H. Vincenti, A. Vernier, F. Quéré, and J. Faure. Vacuum laser acceleration of relativistic electrons using plasma mirror injectors. *Nature Physics*, 12:355 EP –, Dec 2015. Article.
- [18] Andrea Macchi, Marco Borghesi, and Matteo Passoni. Ion acceleration by superintense laser-plasma interaction. *Rev. Mod. Phys.*, 85:751–793, May 2013.
- [19] P. B. Corkum. Plasma perspective on strong field multiphoton ionization. *Phys. Rev. Lett.*, 71:1994–1997, Sep 1993.
- [20] M. Drescher, M. Hentschel, R. Kienberger, M. Uiberacker, V. Yakovlev, A. Scrinzi, Th. Westerwalbesloh, U. Kleineberg, U. Heinzmann, and F. Krausz. Time-resolved atomic inner-shell spectroscopy. *Nature*, 419(6909):803–807, October 2002.
- [21] A. L. Cavalieri, N. Müller, Th. Uphues, V. S. Yakovlev, A. Baltuška, B. Horvath, B. Schmidt, L. Blümel, R. Holzwarth, S. Hendel, M. Drescher, U. Kleineberg, P. M. Echenique, R. Kienberger, F. Krausz, and U. Heinzmann. Attosecond spectroscopy in condensed matter. *Nature*, 449(7165):1029–1032, October 2007.
- [22] M. Uiberacker, Th. Uphues, M. Schultze, A. J. Verhoef, V. Yakovlev, M. F. Kling, J. Rauschenberger, N. M. Kabachnik, H. Schröder, M. Lezius, K. L. Kompa, H.-G. Muller, M. J. J. Vrakking, S. Hendel, U. Kleineberg, U. Heinzmann, M. Drescher, and F. Krausz. Attosecond real-time observation of electron tunnelling in atoms. *Nature*, 446(7136):627–632, April 2007.

-
- [23] Brendan B Godfrey. Numerical cherenkov instabilities in electromagnetic particle codes. *Journal of Computational Physics*, 15(4):504 – 521, 1974.
- [24] G. Blaclard, H. Vincenti, R. Lehe, and J. L. Vay. Pseudospectral maxwell solvers for an accurate modeling of doppler harmonic generation on plasma mirrors with particle-in-cell codes. *Phys. Rev. E*, 96:033305, Sep 2017.
- [25] E-mail: jlvay@lbl.gov Vay, Jean-Luc, Irving Haber, and Brendan B. Godfrey. A domain decomposition method for pseudo-spectral electromagnetic simulations of plasmas. *Journal of Computational Physics*.
- [26] H. Vincenti and F. Quéré. Attosecond lighthouses: How to use spatiotemporally coupled light fields to generate isolated attosecond pulses. *Phys. Rev. Lett.*, 108:113904, Mar 2012.
- [27] A.E. Siegman. *Lasers*. University Science Books, 1986.
- [28] Patrick Mora and R. Pellat. Self-similar expansion of a plasma into a vacuum. *The Physics of Fluids*, 22(12):2300–2304, 1979.
- [29] William L. Kruer and John M. Dawson. The physics of laser plasma interactions. *Physics Today*, 42(8):69–70, August 1989.
- [30] G. Doumy, F. Quéré, O. Gobert, M. Perdrix, Ph. Martin, P. Audebert, J. C. Gauthier, J.-P. Geindre, and T. Wittmann. Complete characterization of a plasma mirror for the production of high-contrast ultraintense laser pulses. *Phys. Rev. E*, 69:026402, Feb 2004.
- [31] B. Dromey, S. Kar, M. Zepf, and P. Foster. The plasma mirror—a subpicosecond optical switch for ultrahigh power lasers. *Review of Scientific Instruments*, 75(3):645–649, 2004.
- [32] T. Wittmann, J. P. Geindre, P. Audebert, R. S. Marjoribanks, J. P. Rousseau, F. Burgy, D. Douillet, T. Lefrou, K. Ta Phuoc, and J. P. Chambaret. Towards ultrahigh-contrast ultraintense laser pulses—complete characterization of a double plasma-mirror pulse cleaner. *Review of Scientific Instruments*, 77(8):083109, 2006.
- [33] Anna Lévy, Tiberio Ceccotti, Pascal D’Oliveira, Fabrice Réau, Michel Perdrix, Fabien Quéré, Pascal Monot, Michel Bougeard, Hervé Lagadec, Philippe Martin, Jean-Paul Geindre, and Patrick Audebert. Double plasma mirror for ultrahigh temporal contrast ultraintense laser pulses. *Opt. Lett.*, 32(3):310–312, Feb 2007.
- [34] L. Chopineau, A. Leblanc, G. Blaclard, A. Denoeud, M. Thévenet, J-L. Vay, G. Bonnaud, Ph. Martin, H. Vincenti, and F. Quéré. Identification of coupling mechanisms between ultraintense laser light and dense plasmas. *Phys. Rev. X*, 9:011050, Mar 2019.
- [35] S. Kahaly, S. Monchocé, H. Vincenti, T. Dzelzainis, B. Dromey, M. Zepf, Ph. Martin, and F. Quéré. Direct observation of density-gradient effects in harmonic generation from plasma mirrors. *Phys. Rev. Lett.*, 110:175001, Apr 2013.

- [36] Alexander Tarasevitch, Konstantin Lobov, Clemens Wünsche, and Dietrich von der Linde. Transition to the relativistic regime in high order harmonic generation. *Phys. Rev. Lett.*, 98:103902, Mar 2007.
- [37] Henri Vincenti. Achieving extreme light intensities using optically curved relativistic plasma mirrors. *Phys. Rev. Lett.*, 123:105001, Sep 2019.
- [38] Gy. Farkas and Cs. Tóth. Proposal for attosecond light pulse generation using laser induced multiple-harmonic conversion processes in rare gases. *Physics Letters A*, 168(5-6):447–450, September 1992.
- [39] A. Bourdier. Oblique incidence of a strong electromagnetic wave on a cold inhomogeneous electron plasma. relativistic effects. *The Physics of Fluids*, 26(7):1804–1807, 1983.
- [40] John David Jackson. *Classical electrodynamics*. Wiley, New York, NY, 3rd ed. edition, 1999.
- [41] F. Quéré, C. Thaury, J-P. Geindre, G. Bonnaud, P. Monot, and Ph. Martin. Phase properties of laser high-order harmonics generated on plasma mirrors. *Phys. Rev. Lett.*, 100:095004, Mar 2008.
- [42] A. Malvache, A. Borot, F. Quéré, and R. Lopez-Martens. Coherent wake emission spectroscopy as a probe of steep plasma density profiles. *Phys. Rev. E*, 87:035101, Mar 2013.
- [43] N. H. Burnett, H. A. Baldis, M. C. Richardson, and G. D. Enright. Harmonic generation in co2 laser target interaction. *Applied Physics Letters*, 31(3):172–174, 1977.
- [44] F. Brunel. Not-so-resonant, resonant absorption. *Phys. Rev. Lett.*, 59:52–55, Jul 1987.
- [45] D. E. Hinkel-Lipsker, B. D. Fried, and G. J. Morales. Analytic expressions for mode conversion in a plasma with a linear density profile. *Physics of Fluids B: Plasma Physics*, 4(3):559–575, March 1992.
- [46] B. Dromey, M. Zepf, A. Gopal, K. Lancaster, M. S. Wei, K. Krushelnick, M. Tatarakis, N. Vakakis, S. Moustazis, R. Kodama, M. Tampo, C. Stoeckl, R. Clarke, H. Habara, D. Neely, S. Karsch, and P. Norreys. High harmonic generation in the relativistic limit. *Nature Physics*, 2(7):456–459, 2006.
- [47] Aleksandar Gjurchinovski. Reflection of light from a uniformly moving mirror. *American Journal of Physics*, 72(10):1316–1324, October 2004.
- [48] H. Vincenti, S. Monchocé, S. Kahaly, G. Bonnaud, Ph Martin, and F. Quéré. Optical properties of relativistic plasma mirrors. *Nature Communications*, 5:3403 EP –, Mar 2014. Article.
- [49] Henri Paul Vincenti. Génération d’impulsions attosecondes sur miroir plasma relativiste. page 224.
- [50] Pedro Arguijo and Marija Strojnik Scholl. Exact ray-trace beam for an off-axis paraboloid surface. *Appl. Opt.*, 42(16):3284–3289, Jun 2003.

- [51] Jonathan A. Wheeler, Antonin Borot, Sylvain Monchocé, Henri Vincenti, Aurélien Ricci, Arnaud Malvache, Rodrigo Lopez-Martens, and Fabien Quéré. Attosecond lighthouses from plasma mirrors. *Nature Photonics*, 6:829 EP –, Nov 2012.
- [52] T. J. Hammond, Graham G. Brown, Kyung Taec Kim, D. M. Villeneuve, and P. B. Corkum. Attosecond pulses measured from the attosecond lighthouse. *Nature Photonics*, 10:171 EP –, Jan 2016.
- [53] Selcuk Akturk, Xun Gu, Pablo Gabolde, and Rick Trebino. The general theory of first-order spatio-temporal distortions of gaussian pulses and beams. *Opt. Express*, 13(21):8642–8661, Oct 2005.
- [54] J. Hebling. Derivation of the pulse front tilt caused by angular dispersion. *Optical and Quantum Electronics*, 28(12):1759–1763, Dec 1996.
- [55] Valentin Gallet. Dispositifs expérimentaux pour la caractérisation spatio-temporelle de chaînes laser femtosecondes haute-puissance. page 170.
- [56] Charles K. Birdsall and A. Bruce Langdon. *Plasma physics via computer simulation*. Taylor and Francis, New York, 2005.
- [57] R. W. Hockney and J. W. Eastwood. *Computer Simulation Using Particles*. Taylor & Francis, Inc., Bristol, PA, USA, 1988.
- [58] J. Derouillat, A. Beck, F. Pérez, T. Vinci, M. Chiaramello, A. Grassi, M. Flé, G. Bouchard, I. Plotnikov, N. Aunai, J. Dargent, C. Riconda, and M. Grech. Smilei : A collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation. *Computer Physics Communications*, 222:351 – 373, 2018.
- [59] J.-L. Vay, A. Almgren, J. Bell, L. Ge, D.P. Grote, M. Hogan, O. Kononenko, R. Lehe, A. Myers, C. Ng, J. Park, R. Ryne, O. Shapoval, M. Thévenet, and W. Zhang. Warp-x: A new exascale computing platform for beam–plasma simulations. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 909:476 – 479, 2018. 3rd European Advanced Accelerator Concepts workshop (EAAC2017).
- [60] Kane Yee. Numerical solution of initial boundary value problems involving maxwell’s equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, May 1966.
- [61] 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.
- [62] Allen Taflove and Susan C. Hagness. *Computational electrodynamics: the finite-difference time-domain method*. Artech House, Norwood, 3rd edition, 2005.
- [63] Henri Vincenti, Mathieu Lobet, Rémi Lehe, R. Sasanka, and Jean-Luc Vay. An efficient and portable SIMD algorithm for charge/current deposition in particle-in-cell codes. *Computer Physics Communications*, 210:145–154, 2017.

- [64] Henri Vincenti and Jean-Luc Vay. Ultrahigh-order maxwell solver with extreme scalability for electromagnetic pic simulations of plasmas. *Computer Physics Communications*, 228:22 – 29, 2018.
- [65] I. Haber, R. Lee, H. Klein, and Boris J. Advances in electromagnetic simulation techniques. In *Proc. Sixth Conf on numerical simulation. Plasmas, Berkeley, CA*, 1973.
- [66] Q. H. Liu. The pseudospectral time-domain (pstd) method: a new algorithm for solutions of maxwell’s equations. In *IEEE Antennas and Propagation Society International Symposium 1997. Digest*, volume 1, pages 122–125 vol.1, July 1997.
- [67] Brendan B. Godfrey. Canonical momenta and numerical instabilities in particle codes. *Journal of Computational Physics*, 19(1):58 – 76, 1975.
- [68] Brendan B. Godfrey, Jean-Luc Vay, and Irving Haber. Numerical stability improvements for the pseudo-spectral em pic algorithm. 2013.
- [69] Brendan B. Godfrey, Jean-Luc Vay, and Irving Haber. Numerical stability analysis of the pseudo-spectral analytical time-domain pic algorithm. *Journal of Computational Physics*, 258:689 – 704, 2014.
- [70] H. Vincenti and J.-L. Vay. Detailed analysis of the effects of stencil spatial variations with arbitrary high-order finite-difference maxwell solver. *Computer Physics Communications*, 200:147 – 167, 2016.
- [71] B. Fornberg. High-order finite differences and the pseudospectral method on staggered grids. *SIAM Journal on Numerical Analysis*, 27(4):904–918, 1990.
- [72] Bengt Fornberg. Generation of finite difference formulas on arbitrarily spaced grids. *Mathematics of Computation*, 51(184):699–699, 1988.
- [73] Jean-Pierre Berenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114(2):185 – 200, 1994.
- [74] Robert I. McLachlan and G. Reinout W. Quispel. Splitting methods. *Acta Numerica*, 11:341–434, 2002.
- [75] Olga Shapoval, Jean-Luc Vay, and Henri Vincenti. Two-step perfectly matched layer for arbitrary-order pseudo-spectral analytical time-domain methods. *Computer Physics Communications*, 235:102 – 110, 2019.
- [76] S. Habib, V. Morozov, H. Finkel, A. Pope, K. Heitmann, K. Kumaran, T. Peterka, J. Insley, D. Daniel, P. Fasel, N. Frontiere, and Z. Lukic. The universe at extreme scale: Multi-petaflop sky simulation on the bg/q. In *SC ’12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pages 1–11, Nov 2012.
- [77] Ian T. Foster and Patrick H. Worley. Parallel algorithms for the spectral transform method. *SIAM J. Sci. Comput.*, 18(3):806–837, May 1997.

- [78] Ananth Grama, editor. *Introduction to Parallel Computing*. Pearson, Harlow, 2. ed., [reprint.] edition, 2011. OCLC: 838852218.
- [79] Amir Gholami, Judith Hill, Dhairya Malhotra, and George Biros. Acefft: A library for distributed-memory FFT on CPU and GPU architectures. *CoRR*, abs/1506.07933, 2015.
- [80] D. Pekurovsky. P3dffft: A framework for parallel computations of fourier transforms in three dimensions. *SIAM Journal on Scientific Computing*, 34(4):C192–C209, 2012.
- [81] Ning Li and Sylvain Laizet. 2DECOMP&FFT - A Highly Scalable 2D Decomposition Library and FFT Interface. page 13.
- [82] M. Frigo and S. G. Johnson. The design and implementation of fftw3. *Proceedings of the IEEE*, 93(2):216–231, Feb 2005.
- [83] Jaewoon Jung, Chigusa Kobayashi, Toshiyuki Imamura, and Yuji Sugita. Parallel implementation of 3d fft with volumetric decomposition schemes for efficient molecular dynamics simulations. *Computer Physics Communications*, 200:57 – 65, 2016.
- [84] S. Jalas, I. Dornmair, R. Lehe, H. Vincenti, J.-L. Vay, M. Kirchen, and A. R. Maier. Accurate modeling of plasma acceleration with arbitrary order pseudo-spectral particle-in-cell methods. *Physics of Plasmas*, 24(3):033115, March 2017.
- [85] Haithem Kallala, Jean-Luc Vay, and Henri Vincenti. A generalized massively parallel ultra-high order fft-based maxwell solver. *Computer Physics Communications*, 244:25 – 34, 2019.
- [86] J.-L. Vay. Noninvariance of space- and time-scale ranges under a lorentz transformation and the implications for the study of relativistic interactions. *Phys. Rev. Lett.*, 98:130405, Mar 2007.
- [87] Peicheng Yu. Lorentz boosted frame simulation technique in Particle-in-cell methods.
- [88] Ferenc Krausz and Misha Ivanov. Attosecond physics. *Rev. Mod. Phys.*, 81:163–234, Feb 2009.
- [89] P. B. Corkum, N. H. Burnett, and M. Y. Ivanov. Subfemtosecond pulses. *Opt. Lett.*, 19(22):1870–1872, Nov 1994.
- [90] G. Sansone, E. Benedetti, F. Calegari, C. Vozzi, L. Avaldi, R. Flammini, L. Poletto, P. Villoresi, C. Altucci, R. Velotta, S. Stagira, S. De Silvestri, and M. Nisoli. Isolated Single-Cycle Attosecond Pulses. 314(5798):443–446.
- [91] F. Ferrari, F. Calegari, M. Lucchini, C. Vozzi, S. Stagira, G. Sansone, and M. Nisoli. High-energy isolated attosecond pulses generated by above-saturation few-cycle fields. *Nature Photonics*, 4:875 EP –, Nov 2010. Article.
- [92] Dmitrii Kormin, Antonin Borot, Guangjin Ma, William Dallari, Boris Bergues, Márk Aladi, István B. Földes, and Laszlo Veisz. Spectral interferometry with waveform-dependent relativistic high-order harmonics from plasma surfaces. *Nature Communications*, 9(1):4992, 2018.

- [93] T. Baeva, S. Gordienko, and A. Pukhov. Relativistic plasma control for single attosecond x-ray burst generation. *Phys. Rev. E*, 74:065401, Dec 2006.
- [94] S G Rykovanov, M Geissler, J Meyer ter Vehn, and G D Tsakiris. Intense single attosecond pulses from surface harmonics using the polarization gating technique. *New Journal of Physics*, 10(2):025025, feb 2008.
- [95] Zi-Yu Chen, Xiao-Ya Li, Bo-Yuan Li, Min Chen, and Feng Liu. Isolated elliptically polarized attosecond soft x-ray with high-brilliance using polarization gating of harmonics from relativistic plasmas at oblique incidence. *Opt. Express*, 26(4):4572–4580, Feb 2018.
- [96] M. Yeung, J. Bierbach, E. Eckner, S. Rykovanov, S. Kuschel, A. Sävert, M. Förster, C. Rödel, G. G. Paulus, S. Cousens, M. Coughlan, B. Dromey, and M. Zepf. Noncollinear polarization gating of attosecond pulse trains in the relativistic regime. *Phys. Rev. Lett.*, 115:193903, Nov 2015.
- [97] Kyung Taec Kim, Chunmei Zhang, Thierry Ruchon, Jean-François Hergott, Thierry Auguste, D. M. Villeneuve, P. B. Corkum, and F. Quéré. Photonic streaking of attosecond pulse trains. *Nature Photonics*, 7(8):651–656, August 2013.
- [98] S. C. Wilks, W. L. Kruer, M. Tabak, and A. B. Langdon. Absorption of ultra-intense laser pulses. *Phys. Rev. Lett.*, 69:1383–1386, Aug 1992.
- [99] B. Dromey, D. Adams, R. Hörlein, Y. Nomura, S. G. Rykovanov, D. C. Carroll, P. S. Foster, S. Kar, K. Markey, P. McKenna, D. Neely, M. Geissler, G. D. Tsakiris, and M. Zepf. Diffraction-limited performance and focusing of high harmonics from relativistic plasmas. *Nature Physics*, 5:146 EP –, Jan 2009. Article.
- [100] J. P. Zou, C. Le Blanc, D. N. Papadopoulos, G. Chériaux, P. Georges, G. Mennerat, F. Druon, L. Lecherbourg, A. Pellegrina, P. Ramirez, F. Giambruno, A. Fréneaux, F. Leconte, D. Badarau, J. M. Boudenne, D. Fournet, T. Valloton, J. L. Paillard, J. L. Veray, M. Pina, P. Monot, J. P. Chambaret, P. Martin, F. Mathieu, P. Audebert, and F. Amiranoff. Design and current progress of the apollon 10 pw project. 3:e2, 2015.
- [101] Sudipta Mondal, Mojtaba Shirozhan, Naveed Ahmed, Maïmouna Bocoum, Frederik Boehle, Aline Vernier, Stefan Haessler, Rodrigo Lopez-Martens, François Sylla, Cedric Sire, Fabien Quéré, Kwinten Nelissen, Katalin Varjú, Dimitris Charalambidis, and Subhendu Kahaly. Surface plasma attosource beamlines at eli-alps. *J. Opt. Soc. Am. B*, 35(5):A93–A102, May 2018.
- [102] Illia Thiele, Stefan Skupin, and Rachel Nuter. Boundary conditions for arbitrarily shaped and tightly focused laser pulses in electromagnetic codes. *Journal of Computational Physics*, 321:1110 – 1119, 2016.
- [103] J. P. Geindre, P. Audebert, and R. S. Marjoribanks. Relativistic *ac* gyromagnetic effects in ultraintense laser-matter interaction. *Phys. Rev. Lett.*, 97:085001, Aug 2006.

- [104] C. Rödel, D. an der Brügge, J. Bierbach, M. Yeung, T. Hahn, B. Dromey, S. Herzer, S. Fuchs, A. Galestian Pour, E. Eckner, M. Behmke, M. Cerchez, O. Jäckel, D. Hemmers, T. Toncian, M. C. Kaluza, A. Belyanin, G. Pretzler, O. Willi, A. Pukhov, M. Zepf, and G. G. Paulus. Harmonic generation from relativistic plasma surfaces in ultrasteep plasma density gradients. *Phys. Rev. Lett.*, 109:125002, Sep 2012.
- [105] Lisandro Dalcin, Mikael Mortensen, and David E Keyes. Fast parallel multidimensional fft using advanced mpi, 2018.
- [106] Pierre Tournois. Acousto-optic programmable dispersive filter for adaptive compensation of group delay time dispersion in laser systems. *Optics Communications*, 140(4):245 – 249, 1997.
- [107] Frédéric Verluise, Vincent Laude, Jean-Pierre Huignard, Pierre Tournois, and Arnold Migus. Arbitrary dispersion control of ultrashort optical pulses with acoustic waves. *J. Opt. Soc. Am. B*, 17(1):138–145, Jan 2000.
- [108] Wolfram Research, Inc. Mathematica, Version 12.0.

Titre: Algorithmes massivement parallèles pour les simulations PIC réalistes de l'interaction laser plasma à ultra-haute intensité, application à la séparation d'impulsions attosecondes d'harmoniques Doppler

Mots clés: Solveurs pseudo-spectraux, FFT en mémoire distribuée, effet phare attoseconde, miroir plasma

Résumé: La complexité des mécanismes physiques mis en jeu lors de l'interaction laser-plasma à ultra-haute intensité nécessite de recourir à des simulations PIC particulièrement lourdes. Au cœur de ces codes de calcul, les solveurs de Maxwell pseudo-spectraux d'ordre élevé présentent de nombreux avantages en termes de précision numérique. Néanmoins, ces solveurs ont un coût élevé en termes de ressources nécessaires. En effet, les techniques de parallélisation existantes pour ces solveurs sont peu performantes au-delà de quelques milliers de coeurs, ou induisent un important usage mémoire, ce qui limite leur scalabilité à large échelle. Dans cette thèse, nous avons développé une toute nouvelle approche de parallélisation qui combine les avantages des méthodes existantes. Cette méthode a été testée à très large échelle et montre un scaling significativement meilleur que les précédentes techniques, tout en garantissant un usage mémoire réduit.

En capitalisant sur ce travail numérique, nous avons réalisé une étude numérique/théorique approfondie dans le cadre de la génération d'harmoniques d'ordres élevés sur cible solide. Lorsqu'une impulsion laser ultra-intense ($I > 10^{16} W.cm^{-2}$) et ultra-courte (de quelques dizaines de femtosecondes) est focalisée sur une cible solide, elle génère un plasma sur-dense, appelé miroir plasma, qui réfléchit non-linéairement le laser incident. La réflexion de l'impulsion laser est accompagnée par l'émission cohérente d'harmoniques d'ordres élevées, sous forme d'impulsions X-UV attosecondes (1 attosecond = $10^{-18}s$). Pour des intensités laser

relativistes ($I > 10^{19} W.cm^{-2}$), la surface du plasma est incurvée sous l'effet de la pression de radiation du laser. De ce fait, les harmoniques rayonnées par la surface du plasma sont focalisées. Dans cette thèse, j'ai étudié la possibilité de produire des impulsions attosecondes isolées en régime relativiste sur miroir plasma, grâce au mécanisme de phare attoseconde. Celui-ci consiste à introduire une rotation des fronts d'onde du laser incident de façon à séparer angulairement les différentes impulsions attosecondes produites à chaque cycle optique. En régime relativiste, la courbure du miroir plasma augmente considérablement la divergence du faisceau harmonique, ce qui rend le mécanisme phare attoseconde inefficace. Pour y remédier, j'ai développé deux techniques de réduction de divergence harmonique afin de mitiger l'effet de focalisation induit par la courbure du miroir plasma et permettre de générer des impulsions attosecondes isolées à partir d'harmoniques Doppler. Ces deux techniques sont basées sur la mise en forme en amplitude et en phase du faisceau laser. Par ailleurs, j'ai développé un modèle théorique pour déterminer les régimes optimaux d'interaction afin de maximiser la séparation angulaire des impulsions attosecondes. Ce modèle a été validé par des simulations numériques PIC en géométries 2D et 3D et sur une large gamme de paramètres laser et plasma. Finalement, on montre qu'en ajustant des paramètres laser et plasma réalistes, il est possible de séparer efficacement les impulsions attosecondes en régime relativiste.

Title: Massively parallel algorithms for realistic PIC simulations of ultra high intensity laser-plasma interaction, application to attosecond pulses separation of Doppler harmonics

Keywords: Pseudo-spectral solvers, distributed memory FFT, attosecond lighthouse effect, plasma mirrors

Abstract: The complexity of the physical mechanisms involved in ultra-high intensity laser-plasma interaction requires the use of particularly heavy PIC simulations. At the heart of these computational codes, high-order pseudo-spectral Maxwell solvers have many advantages in terms of numerical accuracy. This numerical approach comes however with an expensive computational cost. Indeed, existing parallelization methods for pseudo-spectral solvers are only scalable to few tens of thousands of cores, or induce an important memory footprint, which also hinders the scaling of the method at large scales. In this thesis, we developed a novel, arbitrarily scalable, parallelization strategy for pseudo-spectral Maxwell's equations solvers which combines the advantages of existing parallelization techniques. This method proved to be more scalable than previously proposed approaches, while ensuring a significant drop in the total memory use.

By capitalizing on this computational work, we conducted an extensive numerical and theoretical study in the field of high order harmonics generation on solid targets. In this context, when an ultra-intense ($I > 10^{16} W.cm^{-2}$) ultra-short (few tens of femtoseconds) laser pulse irradiates a solid target, a reflective overdense plasma mirror is formed at the target-vacuum interface. The subsequent laser pulse non linear reflection is accompanied with the emission of coherent high order laser harmonics, in the form

of attosecond X-UV light pulses (1 attosecond = $10^{-18}s$). For relativistic laser intensities ($I > 10^{19} W.cm^{-2}$), the plasma surface is curved under the laser radiation pressure. And the plasma mirror acts as a focusing optics for the radiated harmonic beam. In this thesis, we investigated feasible ways for producing isolated attosecond light pulses from relativistic plasma-mirror harmonics, with the so called attosecond lighthouse effect. This effect relies introducing a wavefront rotation on the driving laser pulse in order to send attosecond pulses emitted during different laser optical cycles along different directions. In the case of high order harmonics generated in the relativistic regime, the plasma mirror curvature significantly increases the attosecond pulses divergence and prevents their separation with the attosecond lighthouse scheme. For this matter, we developed two harmonic divergence reduction techniques, based on tailoring the laser pulse phase or amplitude profiles in order to significantly inhibit the plasma mirror focusing effect and allow for a clear separation of attosecond light pulses by reducing the harmonic beam divergence. Furthermore, we developed an analytical model to predict optimal interaction conditions favoring attosecond pulses separation. This model was fully validated with 2D and 3D PIC simulations over a broad range of laser and plasma parameters. In the end, we show that under realistic laser and plasma conditions, it is possible to produce isolated attosecond pulses from Doppler harmonics.

Université Paris-Saclay
Espace Technologique / Immeuble Discovery
Route de l'Orme aux Merisiers RD 128 / 91190 Saint-Aubin, France