University of California
Los Angeles

Lorentz boosted frame simulation technique in Particle-in-cell methods

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

by

Peicheng Yu

2016

Abstract of the Dissertation

Lorentz boosted frame simulation technique in Particle-in-cell methods

by

Peicheng Yu

Doctor of Philosophy in Electrical Engineering

University of California, Los Angeles, 2016

Professor Warren B. Mori, Chair

Accelerators at the energy frontier have been the tool of choice for nearly a century for unraveling the structure of matter, space, and time. Today's accelerators are the most complex and expensive tools for scientific discovery built by humans. The capability of these accelerators has increased at an exponential rate due to the development of new accelerator concepts and technology. The capability of existing accelerator technology has plateaued, so that a future accelerator at the energy frontier will be so large and expensive that it is not clear it will be built. On the other hand, plasma based acceleration has emerged as a possible alternative technology with much recent progress in theory, simulation, and experiment. In plasma based acceleration intense short-pulse laser, or particle beam excites a plasma wave wakefield as it propagates through long regions of plasma. When a laser is used it is called laser wakefield acceleration (LWFA), and when a particle beam is used it is called plasma wakefield acceleration (PWFA). Simulations have contribute greatly to the recent progress by providing guidance and insight for existing experiments, and for permitting the study of parameters beyond the current reach of experiments. However, these simulations require much computing resources. Therefore, alternative numerical techniques are desired, and in some cases are needed.

In this dissertation, we systematically explore the use of a simulation method for modeling LWFA using the particle-in-cell (PIC) method, called the Lorentz boosted frame technique.

In the lab frame the plasma length is typically four orders of magnitude larger than the laser pulse length. Using this technique, simulations are performed in a Lorentz boosted frame in which the plasma length, which is Lorentz contracted, and the laser length, which is Lorentz expanded, are now comparable. This technique has the potential to reduce the computational needs of a LWFA simulation by more than four orders of magnitude, and is useful if there is no or negligible reflection of the laser in the lab frame.

To realize the potential of Lorentz boosted frame simulations for LWFA, the first obstacle to overcome is a robust and violent numerical instability, called the Numerical Cerenkov Instability (NCI), that leads to unphysical energy exchange between relativistically drifting particles and their radiation. This leads to unphysical noise that dwarfs the real physical processes. In this dissertation, we first present a theoretical analysis of this instability, and show that the NCI comes from the unphysical coupling of the electromagnetic (EM) modes and Langmuir modes (both main and aliasing) of the relativistically drifting plasma. We then discuss the methods to eliminate them. In EM-PIC simulations of plasmas, Maxwell's equations are solved using a finite difference form for the derivatives in real space or using FFT's and solving the fields in wave number space. We show that the use of an FFT based solver has useful properties on the location and growth rate of the unstable NCI modes. We first show that the use of an FFT based solver permits the effective elimination of the NCI by both using a low pass filter in wave number space and by reducing the time step. We also show that because some NCI modes are very localized in wave number space, a modification of the numerical dispersion near these unstable modes can eliminate them. We next show that these strategies work just as well if the FFT is only used in the plasma drifting direction and propose a hybrid FFT/Finite Difference solver. This algorithm also includes a correction to the current from the standard charge conserving current deposit that ensures that Gauss's Law is satisfied for the FFT/Finite Difference divergence operator.

However, the use of FFTs can lead to parallel scalability issues when there are many more cells along the drifting direction than in the transverse direction(s). We then describe an algorithm that has the potential to address this issue by using a higher order finite difference operator for the derivative in the plasma drifting direction, while using the standard second

order operators in the transverse direction(s). The NCI for this algorithm is analyzed, and it is shown that the NCI can be eliminated using the same strategies that were used for the hybrid FFT/Finite Difference solver. This scheme also requires a current correction and filtering which require FFTs. However, we show that in this case the FFTs can be done locally on each parallel partition.

We also describe how the use of the hybrid FFT/Finite Difference or the hybrid higher order finite difference/second order finite difference methods permit combining the Lorentz boosted frame simulation technique with another "speed up" technique, called the quasi-3D algorithm, to gain unprecedented speed up for the LWFA simulations. In the quasi-3D algorithm the fields and currents are defined on an $r-z$ PIC grid and expanded in azimuthal harmonics. The expansion is truncated with only a few modes so it has similar computational needs of a 2D $r-z$ PIC code. We show that NCI has similar properties in $r-z$ as in $z-x$ slab geometry and show that the same strategies for eliminating the NCI in Cartesian geometry can be effective for the quasi-3D algorithm leading to the possibility of unprecedented speed up.

We also describe a new code called UPIC-EMMA that is based on fully spectral (FFT) solver. The new code includes implementation of a moving antenna that can launch lasers in the boosted frame. We also describe how the new hybrid algorithms were implemented into OSIRIS. Examples of LWFA using the boosted frame using both UPIC-EMMA and OSIRIS are given, including the comparisons against the lab frame results. We also describe how to efficiently obtain the boosted frame simulations data that are needed to generate the transformed lab frame data, as well as how to use a moving window in the boosted frame.

The NCI is also a major issue for modeling relativistic shocks with PIC algorithm. In relativistic shock simulations two counter-propagating plasmas drifting at relativistic speeds are colliding against each other. We show that the strategies for eliminating the NCI developed in this dissertation are enabling such simulations being run for much longer simulation times, which should open a path for major advances in relativistic shock research.

The dissertation of Peicheng Yu is approved.

Pietro Musumeci

Robert N. Candler

Chandra J. Joshi

Warren B. Mori, Committee Chair

University of California, Los Angeles

2016

To Yaqin, Emma, and Andrew.

To Mom and Dad.

Table of Contents

List of Figures

List of Tables

Acknowledgments

First of all, I would like to thank my Ph.D. advisor, Professor Warren Mori. Having the chance to work with him is simply the luckiest thing in the world. I'll never forget the good old times when we brainstormed together for new ideas to attack a tough physics problem, when we argued over the most accurate words to put in a paper, and when we chatted about interesting things beyond physics.

I would also like to thank my very good friend, and collaborator, Dr. Xinlu Xu. We have been collaborating for over five years, and every day I still learn new knowledges from him. He has countless ways of attacking physics problems. When our interests overlaps, we make progress at the speed of light. And when we work on different aspects of a problem, our work is complementary to each other that adds to the depth of our work.

I am also greatly in debt to Dr. Viktor Decyk. We worked on the production code UPIC-EMMA together. Viktor knows the UPIC framework inside out, and each discussion with him turned out to be very fruitful. It took us less than one year to finish the writing and testing of UPIC-EMMA, which was impossible without his tremendous effort.

A special thanks to Adam Tableman and Asher Davidson, for our collaborations on the quasi-3D Lorentz boosted frame OSIRIS code. This code uses the latest quasi-3D framework that Asher and Adam developed, together with a clean and efficient `fftw` packages wrapped with APIs designed by Adam to be easily fit into OSIRIS. Without their help, this version of the code could not have been accomplished in four month.

I would also like to thank Dr. Weiming An, and his wife Xiaobo Su, for not only the academic collaborations, but also for their friendships and companions for the six years I stayed at UCLA. They are like families to my wife and I.

I would also like to thank all the help from my colleagues Dr. Frank Tsung, and Dr. Ben Winjum. The UCLA plasma simulation group is filled with smart and friendly people, and they made me feel at home during my six-year stay at UCLA.

Vita

2002–2006 B. Eng., Engineering Physics, Tsinghua University, Beijing, China.

2006–2009 M. Eng., Engineering Physics, Tsinghua University, Beijing, China.
– Conducted physics study and engineering design of a laser-electron storage ring for a compact Compton scattering X-ray source.

2009–2010 Researcher, Accelerator Laboratory, Department of Engineering Physics, Tsinghua University, Beijing, China.
– Conducted research on crystalline beam dynamics in ion storage ring.

2010–2016 Ph.D. candidate, Electrical Engineering
University of California, Los Angeles.
– Conducted research on simulating relativistically drifting plasma in Particle-in-cell methods. This includes the theoretical analysis of numerical Cerenkov instability (NCI), the developement of NCI elimination methods in Cartesian and quasi-3D geometry, design of laser wakefield acceleration simulation speedup methods, and implementation of these methods into production codes OSIRIS and UPIC-EMMA.

Publications and Presentations

Xinlu Xu, Peicheng Yu, Samual F. Martins, *et al.*, Comp. Phys. Comm. 184, 2503 (2013).

Peicheng Yu, Xinlu Xu, Viktor K. Decyk, *et al.*, J. Comp. Phys. 266, 124 (2014).

Peicheng Yu, Xinlu Xu, Viktor K. Decyk, *et al.*, Comp. Phys. Comm. 192, 32 (2015).

Peicheng Yu, Xinlu Xu, Adam Tableman, *et al.*, Comp. Phys. Comm. 197, 144 (2015).

Peicheng Yu, Xinlu Xu, Asher Davidson, *et al.*, J. Comp. Phys. 316, 747 (2016)

Fei Li, Peicheng Yu, Xinlu Xu, *et al.*, arXiv:1605.01496 (2016), submitted to Comp. Phys. Comm.

Peicheng Yu, Xinlu Xu, Viktor K. Decyk, *et al.*, Proceedings of Advanced Accelerator Concepts Conference 2012, Austin Texas, AIP Conf. Proc. 1507, 416-420 (2012)

Peicheng Yu, Viktor K. Decyk, Weiming An, *et al.*, Proceedings of Particle Accelerator Conference 2013, Pasadena, CA USA (2013).

Peicheng Yu, Asher Davidson, Adam Tableman, *et al.*, Proceedings of International Particle Accelerator Conference 2015, Richmond, VA, USA (2015).

Peicheng Yu, Asher Davidson, Xinlu Xu, *et al.*, Proceedings of Advanced Accelerator Concepts Conference 2014, San Jose, CA (2014).

Xinlu Xu, Peicheng Yu, Weiming An, *et al.*, Proceedings of Advanced Accelerator Concepts Conference 2012, Austin Texas, AIP Conf. Proc. 1507, 662-665 (2012)

Peicheng Yu, *Numerical instability due to relativistic plasma drift in EM-PIC code*, working group talk at Advanced Accelerator Concepts Conference 2012, Austin, Texas, USA (2012).

Peicheng Yu, *Lorentz boosted frame simulation of Laser wake field acceleration in 3D and quasi-3D geometry*, invited talk at International Conference on Numerical Simulation of Plasmas, Golden, Colorado, USA (2015).

CHAPTER 1

Introduction

Ever since the observation of energetic particles emitted from natural radioactive substances, humans have never stopped their pursuit of finding new particles by colliding them at high energies. Particle accelerators, the machines that produce high energy particle beams, have also seen its own evolution. From electrostatic accelerators that generate particles in the MeV range, to the powerful Large Hadron Collider [1] that delivers a total energy of 13 TeV in the center of mass (with 6.5 TeV in each proton beam), particle accelerators at the energy frontier have led to the discovery of new particles, and to the advancement of science [2]. At the same time, particle accelerators have continually been the most complex and expensive tools built for scientific discovery. High energy colliders [3] at the energy frontier have also produced answers to fundamental questions about the fabric and dynamics of matter, space, and time. High energy particle accelerators are also the key and most expensive component in advanced light sources which produce high energy photon beams which are used to answer fundamental questions in biology, chemistry, and material science.

The output energy of high energy accelerators has continued to increase at an incredible rate owing to the development of new accelerator technologies. From electrostatic accelerators to the linac, and from cyclotrons to storage rings, in the past century accelerator physicists have invented, experimented with, and constructed many different kinds of accelerating structures in order to search for and find new particles. As one technology reached its limit, another technology was invented. In addition, new types of accelerator concepts emerged including colliders [3] where beams were collided against each other in their center of mass, linear colliders, and circular colliders. Present day accelerators use time varying electric fields stored in a phased array of cavities. The electric fields oscillate at radio fre-

quencies, and therefore they are called RF accelerators. As the final energy of the accelerator has increased, so has its overall size and complexity. This has led to the development of secondary technologies such as storage rings, final focuses, and transport lines with thousands of focusing elements. It has also led to the emergence of sub disciplines of physics such as accelerator and beam physics.

There are currently two types of accelerators. Circular and linear accelerators. In a circular accelerator, beams are gradually accelerated as they move in large "circles". The size is determined from the limit of synchrotron radiation. This is why it is preferable to use hadrons because they radiate less at the same energy. In a circular accelerator, beams are accelerated in opposite directions and are collided against each other when the beams reach the desired energy. However, hadrons such as protons (and anti-protons) are not fundamental particles so extra energy is needed to break them apart, and the collision can make a shower of particles. The largest accelerator in the world, the LHC, is 30 kilometers in circumference, and as noted above collides proton beams of 6.5 TeV against each other [1]. There are advantages to colliding leptons, such as electrons and positrons, as they are fundamental particles. However, this requires using a linear collider to avoid energy loss from synchrotron radiation, and the size of a linear collider is limited by the acceleration gradient as the particles cannot recirculate. The largest linear accelerator in the world is still the Stanford Linear Accelerator (SLAC). When operating as a collider it could produce 50 GeV electrons/positrons in 3.2 kilometers. It is now used as the injector of the Linear Coherent Light Source at SLAC [4], which generates coherent photons at x-ray energies. The key acceleration element used in a linac, such as SLAC, is the radio frequency (RF) cavity. The electric field or acceleration gradient in an RF cavity cannot exceed 100 MeV/m, otherwise there will be "RF breakdown" of the cavity walls. This limit on the acceleration gradient has significantly affected the feasibility of building a linear collider at the energy frontier. The high energy physics community has spent several decades on developing a science case and a design to build a new machine called the International Linear Collider (ILC) [5], for the purpose of more refined study of the Higgs Boson and the standard model. In the proposed machine electrons and positrons would collide at 1 TeV in the center of mass, and the length

of the machine would be more than 30 km long based on current RF technology at a cost of more than 20 billion dollars. The huge size and cost of the ILC appears to be too high for even a consortium of countries to build it. Therefore, there is a strong motivation to develop new methods to accelerate particles at much higher acceleration gradients.

There has been a community of researchers working on so-called advanced acceleration concepts since the early 1980s. Among the numerous advanced acceleration concepts that have been proposed over the years, plasma-based acceleration (PBA) has achieved tremendous progress worldwide, and is viewed as the most promising candidate for the basis of building a potentially cheaper linear collider. In 1979, Tajima and Dawson proposed to use an intense short pulse laser to create a wake of plasma oscillations for accelerating electrons [6]. This paper marks the beginning of the field of plasma based acceleration. This was originally called the Wake Plasmon Acceleration, but it is now universally referred to as the Laser WakeField Acceleration (LWFA). In this concept an intense laser pulse propagates in the plasma. The ponderomotive force of the laser pushes the plasma electrons forward and radially away from the ion background. After the laser passes by, the electrons are pulled back by the space charge force from the ion background and oscillate around their equilibrium positions. Such oscillations are phased such that the wake has a phase velocity close to the group velocity of the laser pulse. Since the plasma is already fully broken down there is no breakdown limit. The electric field in the plasma wave wakefield has a longitudinal component, and its amplitude is proportional to the square root of the plasma density. It can reach 100 GeV/m for plasma densities of 10^{18} cm^{-3} which is more than three orders of magnitude higher than in an RF cavity. These three features, a longitudinal component of the electric field, a phase velocity near the speed of light, and no breakdown limit, are ideal for accelerating charged particles to high energy. As a result, LWFA offers the potential to construct compact accelerators that have numerous potential applications, including the building blocks for a next generation linear collider, and being the driver for compact light sources. We note that there is another plasma-based acceleration concept in which the wake is driven by the space charge forces of a relativistic particle beam, and this is called the Plasma WakeField Accelerator (PWFA).

The last decade has seen an incredible amount of theoretical, experimental, and simulation progress in plasma-based acceleration. This has led to the consideration of how to design a future linear collider (LC) using beam or laser drivers [7, 8, 9, 10, 11, 12]. These colliders are based on using at least 20 stages, each one meter long. This progress includes several recent experimental milestones for wakes driven by a charge particle beam, i.e, plasma wakefield acceleration (PWFA). These milestones are the demonstration of sustained high gradient acceleration ($\sim$50GeV/m) over one meter [13], efficient transfer of energy from a drive beam to a trailing beam [14] , and high gradient positron acceleration in self-loaded wakes [15]. Experimental progress in LWFA has also led to several milestones, including the self-injection of electrons into laser-produced wakes [9, 10, 11], the generation of monoenergetic electrons exceeding 1 GeV [12]. Most of this experimental work in PWFA and LWFA has been in the nonlinear blowout regime where the wakefield is multi-dimensional and all (or almost all) of the plasma electrons are expelled forward and radially, leaving behind what is referred to as an ion column, until all of the electrons rush back to create the wake. This nonlinear regime is referred to as the blowout or bubble regime [16, 17, 18].

The physics involved in exciting the wake, and in how the laser or particle beam driver evolves as it propagates until pump depletion is highly nonlinear. It is also difficult to make direct experimental measurements, and therefore it is difficult to interpret the experimental results. As a result, simulations have been very helpful, and in some cases necessary for the rapid progress that has been made in the field of PBA.

When modeling PBA it is necessary to model the evolution of the driver over pump depletion distances, the generation of the wakefield and its evolution, the injection and capture of witness beams of electrons/positrons, and the beam loading of the wakefield by the acceleration of the witness beams. Furthermore, with the current concepts for a linear collider based on PBA relying on 10's of stages, it will be necessary to model how sections are staged. In some schemes such as the afterburner [19] of an existing collider or the use of a compressed proton beam from the LHC, it is also of interest to study PWFA stages that are hundreds of meters rather than a meter long.

Within one PBA stage there are many similarities between modeling PWFA and LWFA.

A key difference is that when modeling a LWFA the smallest spatial scale that needs to be resolved is often the laser wavelength, and not the wavelength of the wakefield. There is often a difference in more than two orders of magnitude between them. In addition, there can be differences between the beam loading scenarios and the spot size of the witness bunch. However, there is much in common in the PIC methods currently used to model the PBA.

1.1 LWFA PIC simulation and Lorentz boosted frame technique

As mentioned above, the physics in LWFA is usually highly nonlinear. As a result, developing predictive theoretical models is challenging [16], and numerical simulations are critical for unrevealing the nonlinear physics LWFA inherents in. The particle-in-cell (PIC) method has been the tool of choice for simulating PBA, including in the seminal paper by Tajima and Dawson [6]. The PIC algorithm follows the self-consistent interactions of particles through the electromagnetic fields directly calculated from the full set of Maxwell equations. Before diving into the details of the modeling of LWFA physics using PIC algorithm, in the following section we will first give a brief introduction of the PIC algorithm commonly used in plasma physics.

1.1.1 Plasma simulation using particle-in-cell method

Ever since the invention of computer, people have been exploring the possibility of modeling plasma physics using computer simulations. The last fifty years has seen numerous techniques for simulating plasma being explored with great success. Among the most successful models for computer simulation of plasmas are particle models [20]. In these models, one simulates the physics by following the motion of a large number of charged particles in their self-consistent electromagnetic fields. While this method sounds straightforward and simple, practical computational limitations require the use of sophisticated techniques [20]. In particular, to significantly reduce the computational loads while preserving the physics, one particle in a simulation can represent many particles of a real plasma and it can be viewed to have a finite size. In the particle-in-cell method, the fields are solved in a discretized

manner using a grid or "cell". Therefore, the fields are defined on grid points. The particle positions have continuous values within the cell determined from the relativistic equation of motion using the Lorentz force. To determine the Lorentz force at the location of the particle, the fields needs to be interpolated to the particle's location based on the values at the grid points. The particles' locations and velocities are used to deposit the charge and current density onto the grid so that Maxwell's equations can be solved to get new fields (Lorentz force) at a later time.

A basic PIC workflow starts with updating the particle positions using the EM fields on the grids. Fields are interpolated onto the particles, and the particles are then pushed by the Lorentz force. After the particle positions and momenta are updated, the values of the source terms, i.e. the current $\vec{j}$ (and charge density ρ for spectral PIC algorithm), are derived from the particle information. The source terms are then plugged into the Maxwell's equation in order to advance the EM fields. These fields are then used to update particles in the next computation cycle.

There are many choices to make when developing an electromagnetic PIC code. One important distinction between codes is how Maxwell's equations are solved. This routine is referred to as the Maxwell solver. There are two types of Maxwell solvers. One is called a finite-difference-time-domain (FDTD) solver [21] and the other is called a spectral solver [20, 22]. The FDTD solver defines the fields on a grid in which the components of the fields are staggered onto grid points shifted by a half-cell size from each other in an appropriate manner on what is called a Yee mesh [21]. This is done so that the derivative of a quantity is defined at the correct location of the grid so that each term in the component of Maxwell's equations is defined at the same location on the grid. In this algorithm it is also common practice to use Ampere's Law and Faraday's Law to advance the EM fields. Gauss' Law is then satisfied through two common methods. In the method most commonly used today, this is done by using a charge conserving current deposition scheme. The charge is never deposited except for use as diagnostic. If Gauss's law is satisfied at the first time step then it is guaranteed to remain satisfied at each ensuing time step. In the other method, the charge density and current density are both deposited, as $\vec{J} = \sum_i q_i \vec{v}_i S(\vec{x}_g - \vec{x}_i)$ and

$Q = \sum_i q_i S(\vec{x}_g - \vec{x}_i)$, where $\vec{x}_g$ are the grid positions $\vec{x}_i$, $\vec{v}_i$ and q_i are the position, velocity, and charge of the i-th particle, and $\vec{S}(\vec{x})$ is the particle shape. Then a correction to the longitudinal component of the electric field is obtained to ensure that Gauss's law is satisfied at each step. This is called the Boris correction. The correction requires a Poisson equation to be solved, which can cause issues when using many computing cores on a parallel computer. The main advantage of a FDTD EM-PIC algorithm with a rigorous charge conserving current deposit [23] is that it can easily be scaled to a large number of nodes because all parts of the algorithm are "local" which means that only information from the nearest cells are needed. It therefore scales well on parallel computers.

In a spectral solver, Maxwell's equations are instead solved in $\vec{k}$ space. The derivatives for each $\vec{k}$ mode are then determined exactly. The electric field is decomposed into longitudinal parts and transverse parts, $\vec{E} = \vec{E}_L + \vec{E}_T$, with $\vec{k} \times \vec{E}_L = 0$ and $\vec{k} \cdot \vec{E}_T = 0$. The longitudinal part of the electric field is obtained at each time step by solving Gauss' Law, while the transverse part is advanced by using Ampere's Law and Faraday's Law. In a spectral EM-PIC algorithm the so-called direct current deposition is used, and there is less numerical noise in its transverse components. The spectral Maxwell solver is generally considered more accurate than a FDTD solver (for all the modes in the fundamental Brillouin zone), and hence this algorithm usually produces cleaner simulation results. The downside of spectral EM-PIC algorithm is that, since the FFT subroutine itself requires global communication, it is more difficult to obtain good parallel scalability. The node that is performing the FFT needs global information in the direction that the FFT is being performed. This requires what is called a transpose of the data. Therefore, to perform an N-dimensional FFT, the problem can only be partitioned in $N-1$ dimensions. After the FFT is performed in the non-partitioned dimension, a transpose is performed so that each node has the global information of the next dimension. To achieve good scalability one has to carefully program the transpose subroutine with complex messaging management, and in some cases adopt complex message passing architecture (e.g. hybrid OPENMP/MPI) in the code. As a result, the development of the spectral Maxwell solver sets a high bar for the computational physicists, and only a few research groups in the world have efficient spectral EM-PIC codes that can scale to

over 100,000 cores. In some cases, it is also more challenging to implement certain types of boundary conditions in a spectral code.

We find that FDTD and spectral EM-PIC codes are often complementary, and when choosing which algorithm to use for a particular problem leads to tradeoffs in accuracy and speed. For some problems, FDTD algorithm is sufficiently accurate, and has greater flexibility for the boundary condition, and 3D domain decomposition can be used. While for other problems, one has to take advantage of the more accurate and isotropic EM dispersion of the spectral solver, and sacrifice some flexibility in parallel partitioning, and computational cost. As we will show in this dissertation, for the elimination of numerical Cerenkov instability (NCI) in multi-dimension, the spectral algorithm gives much cleaner results compared with FDTD algorithm using a Yee mesh. However, it can place certain constraints on the simulation. Hence a large portion of the work described in this dissertation is devoted to determining how to take advantage of both algorithms to achieve the best balance between accuracy and computational cost.

1.1.2 Lorentz boosted frame simulation

When modeling LWFA physics using an EM-PIC code, one has to resolve the smallest physical length of interest, in this case, the laser wavelength which is often more than two orders of magnitude shorter than the wavelength of the wake. For existing short wavelength lasers, $\lambda = 0.8\mu$m, the cell size in the laser propagation direction must be on the order of $\sim 0.1\lambda$. On the other hand, the length of the plasma that the laser driver is propagating through is typically between a centimeter to a meter. The disparity in these two critical physical lengths makes the simulation of LWFA very CPU time consuming, as the algorithm can only advance the laser at most one cell per time step. For instance, using a standard PIC code to study a 10 GeV stage in a nonlinear regime takes approximately 10-100 million core hours (depending on resolution etc.) on today's computers. While computing resources now exist to do a few of 100 million core hour simulations, it is not possible to do parameter scans in full three-dimensions for these problems. Therefore, reduced models such as using

the ponderomotive guiding center approximation (the laser frequency is averaged out of the theory and an envelope equation for the laser and a ponderomotive force term is added to the particle push) with full PIC [24] for the wake or with quasi-static PIC [25, 26] are used for parameter scans. However, while these models are very useful, they cannot as yet model full pump depletion distances and the quasi-static approach cannot as yet model self-injection.

In a 1992 NSF proposal by Mori *et al.*, a novel simulation technique, called the Lorentz boosted frame simulation technique, was proposed [27]. The basic idea was that in a Lorentz boosted frame the plasma length (in the laser propagation direction) is Lorentz contracted while the plasma wake wavelength and laser pulse length are Lorentz expanded. The number of laser cycles is an invariant so the necessary number of cells needed to resolve the laser is also an invariant while the cell size, and hence time step are Lorentz expanded. This idea implicitly assumes that there is negligible, or no reflected laser energy, as the wavelength of the reflected light would be Lorentz contracted and not be properly resolved. The increase in time step and decrease in the plasma length lead to savings of factors of $\gamma_b^2 = (1 - v_b^2/c^2)^{-1}$ as compared to a lab frame simulation using the so-called moving window [28]. In the moving window only the region around the laser is simulated, and the window keeps up with the laser. This proposal was not funded. However, in the early 1990s simulations using the Lorentz boosted frame technique were attempted using the code WAVE. Although unpublished, it was found that in one-dimension that the idea worked. However, when tried in multi-dimensions a robust numerical instability was observed and the work was abandoned.

In 2007, the idea of using a Lorentz boosted frame received a "boost" through the 2007 publication of J-L. Vay [29], who was not aware of the earlier work on this subject. Lorentz boosted frame simulations were carried out using the code OSIRIS [30] as well as WARP [31]. This led to working through ideas on launching a laser as well as how to effectively compare results from the lab frame to the boosted frame. The recent work also found that there was a robust numerical instability for multi-dimensional simulations algorithm. This led to various groups experimenting with strongly filtering the current and the fields. The γ that the plasma drifted across the grid was increased until the instability was observed. It was also found by several groups that the instability was not as strong for an optimum time

step (around $\Delta t = 0.5\Delta x/c$, where $\hat{x}$ is the plasma drifting direction). Despite several years of effort, the root cause of the instability could not be identified.

Little did people realize that the starting point to understanding this instability, now known as the multi-dimensional numerical Cerenkov instability (NCI), was hidden in a 1975 paper by Dr. Brandon Godfrey [35]. However, it was not until nearly forty years later, that people began to unlock the secret of this instability. Interestingly, this instability is also a limiting factor in relativistic shock studies. In these studies the shock is studied by colliding plasmas with $\gamma \sim 20$ or by reflecting a drifting plasma off a wall. While the early work of Godfrey was very useful, the study of NCI in the PIC algorithm is far from complete.

1.2 Numerical Cerenkov Instability: Past and Present

When performing LWFA simulation in the Lorentz boosted frame, an intense laser driver is moving into a counter-propagating plasma. The plasma is drifting relativistically at γ. When performing such a simulation using a PIC code with a standard finite-difference-time-domain (FDTD) Maxwell solver, e.g., Yee solver, a robust numerical instability is observed. The same phenomenon can be observed when there is only a plasma drifting relativistically in the simulation, including in relativistic shock simulations. In the early work on Lorentz boosted frames, the UCLA group attributed the observed instability to an under resolved reflection from of Stimulated Raman Scattering [36]. In the most recent efforts, people began to wonder if it was instead connected to the Numerical Cerenkov Instability (NCI). The numerical instability of a drifting plasma had been previously studied by Dr. Brandon Godfrey which led to two classic papers [34, 35]. In his 1974 paper, Godfrey studied the numerical dispersion of a 1D drifting plasma. He found that for sub-cycling algorithms where the particles are pushed every N time steps of the field solver, that there was an instability if a particle moved more than a cell size during a particle push. He concluded that there was no instability if particles were pushed every time step. In the second paper [35], he studied multi-dimensional effects for thermal plasmas. He found that in multi-dimensions that there was an instability, however, he did not do a relativistically correct analysis. He attributed

the instability to the coupling of light waves to resonant particles through what he referred to as the incompatibility between the Eulerian treatment of the inter-particle forces and the Lagrangian treatment of the particle positions, and through aliasing. Because of the importance of resonant particles the phenomenon is similar to Cerenkov radiation in which a particle can couple to light waves supported by the media with phase velocities less than the speed of light. The difference is that in this case the radiation is slowed down due to numerical errors from the finite grid size and time step applied in the PIC algorithm [37, 38], which leads to the under-resolution of high $|\vec{k}|$ EM modes in the fundamental Brillouin zone, and therefore it was named by Godfrey as the *numerical* Cerenkov instability (NCI).

However, the original work of Godfrey [34, 35] is far from satisfactory, and from what we can tell cannot be directly applied to a relativistically drifting plasma. Specifically, the first paper is 1D in nature, and as noted above is only relevant to sub-cycling algorithms which are not currently used. The second paper, writes out a dispersion relation for the potentials, ϕ and $\vec{A}$, including the effects of a warm plasma. This dispersion relation did not include the relativistic correction, and is therefore not applicable to relativistic drifts. The dispersion relations were not analyzed in detail, and as will be seen in the work presented in this dissertation, is not in a clear form. Therefore, it was not clear if the NCI as described by Godfrey could explain the observations in UCLA's (and LBNL's) early work on the Lorentz boosted frame LWFA simulation, and in the work of relativistic shock simulations.

As noted above it was not until 2012, thirty-seven years after Ref. [35] was published, that the PIC community first started to connect the observed instability in a Lorentz boosted frame to the work in [35]. To our knowledge, the first connection was the work of Yu *et al.*, [39] that was presented at, and published in the proceedings of, the 2012 Advanced Accelerator Concept Workshop at Austin Texas. In this paper, Yu *et al.* showed that the observed instability occurred at the intersection of the particle resonances (fundamental and aliases) and the light wave dispersion relations. A dispersion relation was derived using some of the methods of Godfrey [35]. The authors then began a more careful analysis that properly included relativistic effects as well as that accurately plotted the growth rates for different choices in the field solver and interpolation schemes. The results showed that the

solution to the numerical dispersion relation of a cold relativistically drifting plasma could be used to explain the pattern and growth rates of the instability observed in the PIC simulation involving relativistic plasma drift. The authors showed that an FFT based solver had advantages for potentially mitigating the instability. It is worth noting that, a paper by B. Godfrey and J-L. Vay was submitted on the arXiv (and was subsequently published in Journal of Computational Physics [40]) several days before our journal submittal. This paper [40] analyzed FDTD solvers and predicted the location of the optimal time step. After glancing at this paper, we realized that we missed a term in our analysis that when properly included also predicted that there was an optimal time step (the value depended on the field interpolation scheme) that reduced the growth rate. We updated our work rapidly as most of the content in our paper was not changed. However, as can be seen by reading the two papers, the work is independent. In addition, in our work FFT based solver and NCI in 3D were also considered.

The two groups have continued to work on the NCI. We have shown that the instability is not simply from a coupling between particle resonances and EM modes, but rather from a coupling between two plasma modes. These modes are best identified in the plasma rest frame where one is purely longitudinal (cold plasma or Langmuir oscillation) and the other is purely transverse (a light wave). In the continuum limit these modes cannot couple to each other in any frame. In a new form for the dispersion relation we showed that in the boosted frame there is a coupling term that vanishes in the limit of the time step and cell size approaching zero [46]. The coupling term also depends on the order of the spatial and temporal aliasing. We also showed that it is easiest to mitigate and essentially eliminate the NCI when FFT or higher order solvers are used. Our work and the work of Godfrey, Vay, and co-workers [40, 41, 42, 43, 44] are complementary.

1.3 Outline

In this dissertation, we start in Chapter 2 by presenting a theory that can explain the source of the NCI. Much of Chapter 2 is taken from two published papers [45, 46]. The theory

also provides a way to accurately and rapidly calculate the patterns and growth rates of these NCI modes. In [46], we showed that the NCI can be explained as the unphysical coupling between the EM modes and Langmuir modes (both main and aliasing), which is fundamentally inevitable due to the fact that we are using finite grid sizes and finite time steps in the PIC system.

Based on this theory, in Chapter 3 we developed strategies for systematically eliminating the unphysical modes. These include the use of a spectral Maxwell solver that spatially advances Maxwell equations in Fourier space, and the modification of the EM dispersion near the location of the unphysical coupling between the EM and Langmuir modes. In Chapter 4, we exploit the insight from Chapter 3 to develop a hybrid Yee-FFT solver that has similar NCI properties as the fully spectral solver, but which has some advantages in parallel scalability, and in boundary conditions along the directions that are not Fourier transformed. This idea is then extended to a customized Maxwell solver in Chapter 5, in which we apply a customized spatial derivative in the plasma drifting direction. This scheme corrects the current in $\vec{k}$ space in the drifting direction so that Gauss' Law is satisfied. This is done on each local domain so that global communication in the FFT is not required. The customized solver allows the use of completely FFT-free Maxwell solvers (only 1D local FFT is needed for the current) that has excellent scalability in all directions, and which makes it possible to apply the NCI mitigation schemes for FFT solvers. The idea has been tested in both LWFA, and relativistic collisionless shock simulations [47, 48, 49, 50].

The fact that one only needs to change the spatial derivatives in the plasma drifting direction for the NCI elimination indicates that the hybrid Yee-FFT solver, and customized solver methods can both be extended to cylindrical/quasi-3D geometries, as will be discussed Chapter 4, 5, and 7. The quasi-3D algorithm is another reduced model that has been recently proposed to speedup LWFA simulations [51]. The idea is to treat the numerical system as PIC in $r - z$, and gridless in azimuthal angle ϕ. The electromagnetic fields and currents are decomposed and truncated into azimuthal modes, and the Maxwell equations are then solved mode by mode. The corresponding EM fields are subsequently added up for updating the particle positions and momenta. The particles are defined in a full 3D geometry, but the

current is deposited only on the $r-z$ grid. As a result, this algorithm can reduce modeling a 3D problem with low azimuthal asymmetry into the similar computational cost as using a 2D $r-z$ code. In Chapter 7, we apply the NCI elimination into the quasi-3D geometry, and achieve the combination of Lorentz boosted frame technique and the quasi-3D technique. This results in an unprecedented speed up of LWFA simulations.

In Chapter 6 and 7, we will also describe details for how we modified our production codes, UPIC-EMMA and OSIRIS, to eliminate the NCI, and present additional sample simulation results of LWFA in both Cartesian and quasi-3D geometry. In Chapter 6 and 7, details regarding the simulation setups of performing Lorentz boosted frame simulation will likewise be presented. This includes the output and post-processing of boosted frame simulation datas to conveniently transform them back to the lab frame, and the choice of moving window in the boosted frame. Directions for future work are presented in Chapter 8.

In the Appendices, we define notation for the numerical dispersion of drifting plasmas, and present the application of NCI elimination scheme in other simulation scenarios, e.g. down ramp injection simulation for FEL injector (Appendix B). We show that by applying the NCI elimination schemes in the lab frame PIC simulation of plasma-based accelerator, we can eliminate the unphysical emittance growth imposed by the grid, and obtain a more accurate estimate of injected beam emittance.

CHAPTER 2

Theory of numerical Cerenkov instability

As mentioned in Chapter 1, in a LWFA simulation the smallest length of physical interests is the wavelength of the driver laser. In the corresponding PIC simulation, one needs to resolve this wavelength, which is on the scale of 1μm. On the other hand, the plasma column that the laser driver propagates through is on the scale of a meter. In the LWFA PIC simulation the laser is advanced a distance less than the smallest cell size in each time step. Due to the disparity between the cell size and the length of the plasma column, LWFA PIC simulations require large number of time steps, and therefore this kind of simulations become very CPU-time consuming. To resolve this disparity in the two physical lengths, one can perform the simulation in a so-called Lorentz boosted frame [27, 29]. In this frame, the laser wavelength will be stretched by $\gamma(1+\beta)$, while the plasma column will be contracted by γ, where γ is the Lorentz factor of the boosted frame. As a result, one can use a much larger cell size and hence much larger time step, and use a plasma column that is much shorter. Assuming that the number of particles per cell remain fixed, this will lead to a theoretical speed up on the scale of $\sim \gamma^2$ as compared to lab frame simulations that were already using the moving window.

However, in the Lorentz boosted frame we are simulating the head-on collision between a laser pulse, and a plasma that is drifting *relativistically* across the grid. As mentioned in Chapter 1, when the plasma is drifting relativistically across the grid, a violent instability will rapidly grow in the simulation. One can observe this numerical artifact by filling a 2D simulation box with plasma drifting relativistically in the x_1 direction. Apply periodic boundary condition in both the x_1 and x_2 directions, and give the plasma a very small temperature to seed the instability. By tracking the total EM energy in the box, it is clear

that there is an exponential growth of the EM energy as the simulation proceeds, and the growth rate of this instability is high. A drifting plasma should be stable in reality, and this exponential growth of EM energy is clearly a *numerical* artifact.

In the following sections we will start with a theoretical analysis of the simple scenario of a 2D drifting plasma as described above. We derive the numerical dispersion relation of this drifting plasma. We will discuss the exact form of the dispersion when the distribution function of a cold relativistically drifting plasma is applied. After that, we will discuss our approaches to solve the dispersion relation, and obtain the expressions of the NCI modes. As we can see later in this chapter, examining the dispersion relation will shed light on the fundamental cause of this instability, and will also give us directions on how to eliminate them.

2.1 Derivation of the numerical dispersion relation

We will follow the notation in Ref. [35] to derive the numerical dispersion relation for a cold plasma drifting with relativistic velocities. The multi-dimensional analysis in [35] solves for scalar potential ϕ and vector potential $\vec{A}$ and is not valid for relativistic drifts. In our analysis we include relativistic mass effects, and provide a framework for studying the effect of using different types of field solvers and current deposition schemes. Since most EM-PIC codes now in use solve for the electric field $\vec{E}$ and magnetic field $\vec{B}$ directly (with finite difference or spectral solvers), we derive a numerical dispersion relation directly using these two quantities. Gaussian units will be used; in addition, particle mass and velocity will be normalized to electron mass and the speed of light.

For a multi-dimensional simulation setup in Cartesian coordinates, the EM field that is interpolated on a particle can be expressed as

$$\begin{aligned}\vec{E}(t,\vec{x}) &= \sum_{m,\vec{n}} \overleftrightarrow{S_E}(t,m,\vec{x},\vec{n})\vec{E}_{m,\vec{n}} \\ \vec{B}(t,\vec{x}) &= \sum_{m,\vec{n}} \overleftrightarrow{S_B}(t,m,\vec{x},\vec{n})\vec{B}_{m,\vec{n}}\end{aligned} \tag{2.1}$$

where m is the time index and $\vec{n}$ is the grid index; $\overleftrightarrow{S}$ are the interpolation tensors used to obtain the appropriate field at $\vec{x}$ and $t = m\Delta t$; $\vec{E}_{m,\vec{n}}$ and $\vec{B}_{m,\vec{n}}$ stands for the electromagnetic forces at time grid index m and space grid index $\vec{n}$. For a momentum conserving field interpolation method $\overleftrightarrow{S_E}$ and $\overleftrightarrow{S_B}$ are equal and are scalar functions times the unit tensor while for the energy conserving field interpolation method $\overleftrightarrow{S_E}$ and $\overleftrightarrow{S_B}$ are not equal in each direction. We have listed typical interpolation functions used in a EM-PIC code in Appendix A.2. The momentum change of the particle is related to the change in the distribution function of the plasma by the linearized Vlasov equation

$$\frac{\partial}{\partial t} f(t,\vec{x},\vec{p}) + \frac{\vec{p}}{\gamma}\cdot\frac{\partial}{\partial \vec{x}} f(t,\vec{x},\vec{p}) + q\left\{\vec{E}(t,\vec{x}) + \frac{\vec{p}}{\gamma}\times\vec{B}(t,\vec{x})\right\}\cdot\frac{\partial f_0}{\partial \vec{p}} = 0$$

where $\vec{p}$ is the particle momentum, and γ is the particle Lorentz factor. After Fourier transforming, the Vlasov Equation becomes

$$f(\omega,\vec{k},\vec{p}) = -iq\left\{\overleftrightarrow{S_E}(\omega,\vec{k})\vec{E}(\omega,\vec{k}) + \frac{\vec{p}}{\gamma}\times\{\overleftrightarrow{S_B}(\omega,\vec{k})\vec{B}(\omega,\vec{k})\}\right\}\cdot\frac{\partial f_0}{\partial \vec{p}}(\omega - \vec{k}\cdot\frac{\vec{p}}{\gamma})^{-1} \qquad (2.2)$$

Note that $\vec{E}$ and $\vec{B}$ are defined at discrete grid positions and discrete values of time, so its Fourier transform in $(w,\vec{k})$ is periodic, i.e.,

$$\vec{E}(\omega,\vec{k}) = \vec{E}(\omega',\vec{k}') \qquad \vec{B}(\omega,\vec{k}) = \vec{B}(\omega',\vec{k}') \qquad (2.3)$$

where

$$\begin{aligned} &\omega' = \omega + \mu\omega_g \qquad \omega_g = \frac{2\pi}{\Delta t} \qquad \mu = 0,\pm1,\pm2,\ldots \\ &k_i' = k_i + \nu_i k_{gi} \qquad k_{gi} = \frac{2\pi}{\Delta x_i} \qquad \nu_i = 0,\pm1,\pm2,\ldots \end{aligned} \qquad (2.4)$$

Note that when the EM fields are staggered (such as on a Yee mesh), there is an additional $(-1)^{\sum_i \nu_i}$ term for each component $\vec{E}(\omega',\vec{k}')$ and $\vec{B}(\omega',\vec{k}')$, where $\hat{i}$ is summed over the directions for which the specific component of the EM field is staggered a half-grid offset from where charge density is defined. Details for where this term arises from can be found in Appendix A.1. We absorb these additional coefficients into the quantities $\overleftrightarrow{S_E}$ and $\overleftrightarrow{S_B}$ to keep Eq. (2.1) correct when $\vec{n}$ includes only integer indices.

Replacing $(w, \vec{k})$ with $(w', \vec{k}')$ in Eq. (2.2), and using Eq. (2.4), we obtain

$$f(\omega', \vec{k}', \vec{p}) = -iq\Big\{\overleftrightarrow{S_E}(\omega', \vec{k}')\vec{E}(\omega, \vec{k}) + \frac{\vec{p}}{\gamma} \times \{\overleftrightarrow{S_B}(\omega', \vec{k}')\vec{B}(\omega, \vec{k})\}\Big\} \cdot \frac{\partial f_0}{\partial \vec{p}}(\omega' - \vec{k}' \cdot \frac{\vec{p}}{\gamma})^{-1} \tag{2.5}$$

The current density $\vec{j}$ due to the movement of the particles can be expressed as

$$\vec{j}(t, \vec{x}) = q \int \overleftrightarrow{S_j}(\vec{x'} - \vec{x})\frac{\vec{p}}{\gamma} f(\{m + 1/2\}\Delta t, \vec{x}', \vec{p}) d\vec{x}' d\vec{p} \tag{2.6}$$

where $\overleftrightarrow{S_j}(\vec{x}' - \vec{x})$ is the tensor for the current deposit. After Fourier transforming we obtain

$$\vec{j}(\omega, \vec{k}) = q \sum_{\mu, \vec{\nu}} (-1)^{\mu} \int d\vec{p} \frac{\overleftrightarrow{S_j}(-\vec{k}')\vec{p}}{\gamma} f(\omega', \vec{k}', \vec{p}) \tag{2.7}$$

We can now proceed in the normal way to obtain a dispersion relation. We start from Faraday's and Ampere's Law,

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t}$$
$$\nabla \times \vec{B} = \frac{\partial \vec{E}}{\partial t} + 4\pi\vec{j}$$

which upon Fourier transforming gives,

$$[\vec{k}]_E \times \vec{E} = [\omega]\vec{B} \tag{2.8}$$

$$[\vec{k}]_B \times \vec{B} = -[\omega]\vec{E} - 4\pi i\vec{j} \tag{2.9}$$

where $[k]_E$ and $[k]_B$ are the finite difference operators for the specific Maxwell solver schemes being used to solve for the $\vec{E}$ and $\vec{B}$ fields. We follow the notation in Ref. [35], and use $[\,\cdot\,]$ exclusively to indicate the Fourier representation for the finite difference operator. Applying $[\vec{k}]_B\times$ to both sides of Eq. (2.8), and using Eq. (2.9), we end up with the coupled wave equation for $\vec{E}$ and $\vec{j}$,

$$([\omega]^2 - [\vec{k}]_E \cdot [\vec{k}]_B + [\vec{k}]_E[\vec{k}]_B)\vec{E} = -4\pi i[\omega]\vec{j} \tag{2.10}$$

Using Eq. (2.7) and (2.10), we could obtain

$$([\omega]^2 - [\vec{k}]_E \cdot [\vec{k}]_B + [\vec{k}]_E[\vec{k}]_B)\vec{E} = -4\pi iq \sum_{\mu, \vec{\nu}} (-1)^{\mu}[\omega] \int d\vec{p} \frac{\overleftrightarrow{S_j}(-\vec{k}')\vec{p}}{\gamma} f(\omega', \vec{k}', \vec{p}) \tag{2.11}$$

If we normalize the distribution function such that $f_0 = n_0 f_0^n$, use the definition of plasma frequency

$$\omega_p^2 = \frac{4\pi q^2 n_0}{m} \tag{2.12}$$

and use the expression for the distribution function in Eq. (2.5), we finally obtain [39]

$$\begin{aligned}
&([\omega]^2 - [\vec{k}]_E \cdot [\vec{k}]_B + [\vec{k}]_E [\vec{k}]_B)\vec{E} \\
= \; & -\omega_p^2 \sum_{\mu,\vec{\nu}} (-1)^\mu \left\{ \int \frac{\overleftrightarrow{S_j}(-\vec{k}')\vec{p}d\vec{p}}{\gamma\omega' - \vec{k}'\cdot\vec{p}} \left\{ [\omega]\overleftrightarrow{S_E}(\omega', \vec{k}')\vec{E} + \frac{\vec{p}}{\gamma} \times \{\overleftrightarrow{S_B}(\omega', \vec{k}')([\vec{k}]_E \times \vec{E})\} \right\} \cdot \frac{\partial f_0}{\partial \vec{p}} \right\}
\end{aligned} \tag{2.13}$$

which is a generalized dispersion relation for a plasma of finite size particles drifting on a grid. We note that the use of additional smoothers and filters can be incorporated into the dispersion relation by adding additional $S_{SM}(\vec{k}')$ terms outside the summation over Brillouin zones (essentially it multiplies the ω_p^2 term).

We next examine the dispersion relation in the limit of a cold plasma including the possibility that the drift is near the speed of light. Note that $\overleftrightarrow{S}$ for the fields and current has only three diagonal elements S_1, S_2, S_3 in each case. In 3D, we can expand Eq. (2.13) explicitly as

$$\begin{aligned}
&\begin{pmatrix}
([\omega]^2 - \sum_{i=1}^3 [k]_{Ei}[k]_{Bi})E_1 + [k]_{E1}[k]_{B1}E_1 + [k]_{E1}[k]_{B2}E_2 + [k]_{E1}[k]_{B3}E_3 \\
([\omega]^2 - \sum_{i=1}^3 [k]_{Ei}[k]_{Bi})E_2 + [k]_{E2}[k]_{B1}E_1 + [k]_{E2}[k]_{B2}E_2 + [k]_{E2}[k]_{B3}E_3 \\
([\omega]^2 - \sum_{i=1}^3 [k]_{Ei}[k]_{Bi})E_3 + [k]_{E3}[k]_{B1}E_1 + [k]_{E3}[k]_{B2}E_2 + [k]_{E3}[k]_{B3}E_3
\end{pmatrix} \\
= \; & -\omega_p^2 \sum_{\mu,\vec{\nu}} (-1)^\mu \int \frac{dp_1 dp_2 dp_3}{\gamma(\gamma\omega' - k_1' p_1 - k_2' p_2 - k_3' p_3)} \begin{pmatrix} S_{j1}p_1 \\ S_{j2}p_2 \\ S_{j3}p_3 \end{pmatrix} \begin{pmatrix} \partial f_0^n/\partial p_1 \\ \partial f_0^n/\partial p_2 \\ \partial f_0^n/\partial p_3 \end{pmatrix}^T \cdot \\
&\begin{pmatrix}
\gamma[\omega]S_{E1}E_1 + p_2 S_{B3}([k]_{E1}E_2 - [k]_{E2}E_1) + p_3 S_{B2}([k]_{E1}E_3 - [k]_{E3}E_1) \\
\gamma[\omega]S_{E2}E_2 + p_3 S_{B1}([k]_{E2}E_3 - [k]_{E3}E_2) + p_1 S_{B3}([k]_{E2}E_1 - [k]_{E1}E_2) \\
\gamma[\omega]S_{E3}E_3 + p_1 S_{B2}([k]_{E3}E_1 - [k]_{E1}E_3) + p_2 S_{B1}([k]_{E3}E_2 - [k]_{E2}E_3)
\end{pmatrix}
\end{aligned} \tag{2.14}$$

This can be rewritten as

$$\overleftrightarrow{\epsilon}(\omega,k)\vec{E} = \begin{pmatrix} \epsilon_{11} & \epsilon_{12} & \epsilon_{13} \\ \epsilon_{21} & \epsilon_{22} & \epsilon_{23} \\ \epsilon_{31} & \epsilon_{32} & \epsilon_{33} \end{pmatrix} \begin{pmatrix} E_1 \\ E_2 \\ E_3 \end{pmatrix} = 0 \tag{2.15}$$

where we note that $\overleftrightarrow{\epsilon}$ is not the dielectric tensor. In addition, we are most interested in a cold plasma that is drifting relativistically. For such a case, the unperturbed normalized distribution function is given by

$$f_0^n = \delta(p_1 - p_0)\delta(p_2)\delta(p_3) \tag{2.16}$$

where $p_0 = \gamma v_0$, and v_0 is the drifting velocity of the plasma. Substituting the above form for f_0^n, Eq. (2.16), into Eq. (2.14), and carrying out the integration we obtain all the elements in the tensor as

$$\begin{aligned}
\epsilon_{11} &= [\omega]^2 - [k]_{E2}[k]_{B2} - [k]_{E3}[k]_{B3} \\
&\quad - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^\mu \frac{S_{j1}\{S_{E1}[\omega]\omega'/\gamma^2 + v_0^2(S_{B3}k_2'[k]_{E2} + S_{B2}k_3'[k]_{E3})\}}{(\omega' - k_1'v_0)^2} \\
\epsilon_{12} &= [k]_{E1}[k]_{B2} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^\mu \frac{S_{j1}v_0k_2'(S_{E2}[\omega] - v_0S_{B3}[k]_{E1})}{(\omega' - k_1'v_0)^2} \\
\epsilon_{13} &= [k]_{E1}[k]_{B3} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^\mu \frac{S_{j1}v_0k_3'(S_{E3}[\omega] - v_0S_{B2}[k]_{E1})}{(\omega' - k_1'v_0)^2} \\
\epsilon_{21} &= [k]_{E2}[k]_{B1} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^\mu \frac{v_0S_{j2}S_{B3}[k]_{E2}}{\omega' - k_1'v_0} \\
\epsilon_{22} &= [\omega]^2 - [k]_{E1}[k]_{B1} - [k]_{E3}[k]_{B3} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^\mu \frac{S_{j2}(S_{E2}[\omega] - v_0S_{B3}[k]_{E1})}{\omega' - k_1'v_0} \\
\epsilon_{23} &= [k]_{E2}[k]_{B3} \\
\epsilon_{31} &= [k]_{E3}[k]_{B1} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^\mu \frac{v_0S_{j3}S_{B2}[k]_{E3}}{\omega' - k_1'v_0} \\
\epsilon_{32} &= [k]_{E3}[k]_{B2} \\
\epsilon_{33} &= [\omega]^2 - [k]_{E1}[k]_{B1} - [k]_{E2}[k]_{B2} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^\mu \frac{S_{j3}(S_{E3}[\omega] - v_0S_{B2}[k]_{E1})}{\omega' - k_1'v_0}
\end{aligned} \tag{2.17}$$

The dispersion relation is then finally obtained from the condition that

$$\mathrm{Det}(\overleftrightarrow{\epsilon}) = 0 \tag{2.18}$$

which is valid in any number of dimensions.

Much can be learned from examining the 1D and 2D limits to the general dispersion relation. In 1D simulations all physical quantities only depend on one coordinate x_1, hence $[\vec{k}]$, $\vec{k}$, and $\vec{k}'$ only have the $\hat{1}$-component. It follows then the elements of $\overleftrightarrow{\epsilon}$ are

$$\begin{aligned}
\epsilon_{11} &= [\omega]^2 - \frac{\omega_p^2}{\gamma}\sum_{\mu,\nu}(-1)^{\mu}\frac{S_{j1}S_{E1}[\omega]\omega'/\gamma^2}{(\omega' - k_1' v_0)^2} \\
\epsilon_{22} &= [\omega]^2 - [k]_{E1}[k]_{B1} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\nu}(-1)^{\mu}\frac{S_{j2}(S_{E2}[\omega] - S_{B3}[k]_{E1}v_0)}{\omega' - k_1' v_0} \\
\epsilon_{33} &= [\omega]^2 - [k]_{E1}[k]_{B1} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\nu}(-1)^{\mu}\frac{S_{j3}(S_{E3}[\omega] - S_{B2}[k]_{E1}v_0)}{\omega' - k_1' v_0} \\
\epsilon_{12} &= \epsilon_{13} = \epsilon_{21} = \epsilon_{23} = \epsilon_{31} = \epsilon_{32} = 0
\end{aligned} \tag{2.19}$$

Using Eq. (2.18), the dispersion relation for the 1D case consists of three uncoupled modes,

$$\epsilon_{11} = 0 \qquad \epsilon_{22} = 0 \qquad \epsilon_{33} = 0 \tag{2.20}$$

where each mode corresponds to separate components of the electric fields E_1, E_2, and E_3 respectively. Each of these modes is numerically stable as long as Δt is sufficiently small. If we take the limit $\Delta t \to 0$, and $\Delta x \to 0$, then Eq. (2.19) and (2.20) reduce to the dispersion relations in a real drifting plasma (which is completely stable). Therefore, there is no NCI in one dimension, which is consistent with the early work of Godfrey [34].

Similarly, the elements of $\overleftrightarrow{\epsilon}$ in the 2D limit can be written as

$$\epsilon_{11} = [\omega]^2 - [k]_{E2}[k]_{B2} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^{\mu}\frac{S_{j1}(S_{E1}\omega'[\omega]/\gamma^2 + S_{B3}[k]_{E2}k_2'v_0^2)}{(\omega' - k_1'v_0)^2}$$

$$\epsilon_{12} = [k]_{E1}[k]_{B2} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^{\mu}\frac{k_2'v_0S_{j1}(S_{E2}[\omega] - S_{B3}v_0[k]_{E1})}{(\omega' - k_1'v_0)^2}$$

$$\epsilon_{21} = [k]_{E2}[k]_{B1} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^{\mu}\frac{S_{j2}S_{B3}[k]_{E2}v_0}{\omega' - k_1'v_0}$$

$$\epsilon_{22} = [\omega]^2 - [k]_{E1}[k]_{B1} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^{\mu}\frac{S_{j2}(S_{E2}[\omega] - S_{B3}[k]_{E1}v_0)}{\omega' - k_1'v_0}$$

$$\epsilon_{33} = [\omega]^2 - [k]_{E1}[k]_{B1} - [k]_{E2}[k]_{B2} - \frac{\omega_p^2}{\gamma}\sum_{\mu,\vec{\nu}}(-1)^{\mu}\frac{S_{j3}(S_{E3}[\omega] - S_{B2}[k]_{E1}v_0)}{\omega' - k_1'v_0}$$

$$\epsilon_{13} = \epsilon_{23} = \epsilon_{31} = \epsilon_{32} = 0 \tag{2.21}$$

Using Eq. (2.18), we can obtain the dispersion relation for the 2D case

$$\epsilon_{11}\epsilon_{22} - \epsilon_{12}\epsilon_{21} = 0 \qquad \epsilon_{33} = 0 \tag{2.22}$$

Note that E_3 is de-coupled from the other two directions.

2.2 Numerical solution v.s. simulation

The key to understanding the NCI from the numerical dispersion relation of drifting plasma is to solve Eq. (2.17) and (2.18). In the following sections, we will solve it both numerically and analytically. For the numerical method, by plugging the $[\vec{k}]$ for a particular Maxwell solver, and numerically solving Eq. (2.18) and looking for the unstable modes, one can obtain the patterns and growth rates of NCI for that Maxwell solver. Correspondingly, one can also directly perform PIC simulations to observe and measure the NCI by applying the same parameters used in the numerical evaluation, and compare the simulation results against the numerical results.

2.2.1 Numerical solution

Without loss of generality, we study the numerical instability induced by the relativistic plasma drift in a 2D system. According to the dispersion relation in 2D, we expect to observe instability in E_2 (and B_3). By calculating the maximum imaginary part of ω for real values of (k_1, k_2) for Eq. (2.22), we can obtain the characteristic pattern of the instability in Fourier space, as well as the growth rate of the instability. We can also plot the real part of ω for k_1, k_2. These results can be used later to compare with the simulation results.

The dispersion relation is general and can be used to examine how the NCI depends on different choices for the Maxwell solvers, on different choices on the field interpolation, e.g. energy and momentum conserving field interpolation, and on the use of smoothing and low pass filters. In this chapter, we show that our dispersion relation agrees well with the simulation results for the cases studied; that we can predict the region of unstable modes by plotting where the beam modes (resonances) and EM modes intersect in $\vec{k}$ and ω space; that we can obtain an asymptotic expression for the growth in 3D which agrees well with the simulations for various finite difference solvers (including the optimal time step that minimize the growth rate); and there are advantages for using a spectral (FFT-based) solver from the point of view of eliminating the instability. Nonetheless, we do not attempt to carry out a comprehensive survey of all available choices listed above in this chapter.

We illustrate the instability using a 2D case with the standard Yee solver [21]. We choose the grid parameters and time step that satisfies the Courant Condition to eliminate the well known numerical instability for the EM modes. We use the parameters in Tab. 2.2, and substitute the finite difference operators for the Yee solver into the 2D dispersion relation. We assume linear (area) interpolation, momentum conserving field interpolation, and a charge conserving current deposition.

After obtaining all the roots (ω, k_1, k_2), we plot the dependence of the growth rate in the (ω_r, k_1) space [figure 2.1 (c)], as well as in the (k_1, k_2) space [figure 2.1 (d)]. It is evident that all the instabilities are near the main or aliased beam modes (resonances). Since the terms with $|\mu| \leq 1$, and $|\nu_1| \leq 1$ are the most important, we neglect higher order terms when

solving Eq. (2.22). Higher order μ and $\vec{\nu}$ terms can be included in the summation if needed. These additional terms lead to additional unstable modes in (k_1, k_2) space with lower growth rates as well as to small modifications to the growth rate and location of the original modes. A plot in (k_1, k_2) space with more terms included are presented in figure 2.1 (b).

While the results in figure 2.1 (c) and (d) are numerically calculated from Eq. (2.22), the location of the unstable modes can also be conveniently predicted by plotting the intersection of the EM modes and beam resonances in (k_1, k_2, ω_r) space. This is shown in a 3D plot [figure 2.1 (a)]. By examining the unstable pattern in (k_1, k_2) space we see that the central part of the pattern comes from the intersections of the EM modes and main beam resonance ($\mu = 0$ and $\nu = 0$), while the part at the four corners can be identified with the intersections of the EM modes and first order spatial aliasing beam resonances ($\mu = 0$ and $\nu_1 = \pm 1$). Making a plot in (k_1, k_2, ω_r) of the intersection of the EM modes and beam resonances for various solvers becomes a useful method for examining where the unstable modes reside without having to solve the full dispersion relation. We note that although this permits a quick survey of the location of the unstable modes, it does not mean the instability arises from the coupling between EM modes and "beam resonance". As we show later, the NCI modes arises from the coupling between modes which are purely transverse (EM) and longitudinal (Langmuir) in the rest frame of the plasma.

Parameters	**Values**
solver	Yee
grid size $(k_p\Delta x_1, k_p\Delta x_2)$	$(0.1, 0.1)$
time step $\omega_p\Delta t$	0.9×Courant limit
boundary condition	Periodic
simulation box size (k_pL_1, k_pL_2)	51.2×25.6
plasma drifting Lorentz factor	$\gamma = 50.0$
plasma density	$n/n_p = 1$

Table 2.1: Simulation parameters for the 2D relativistic plasma drift simulation. In the table $k_p = \omega_p/c$, where ω_p is defined in Eq. (2.12).

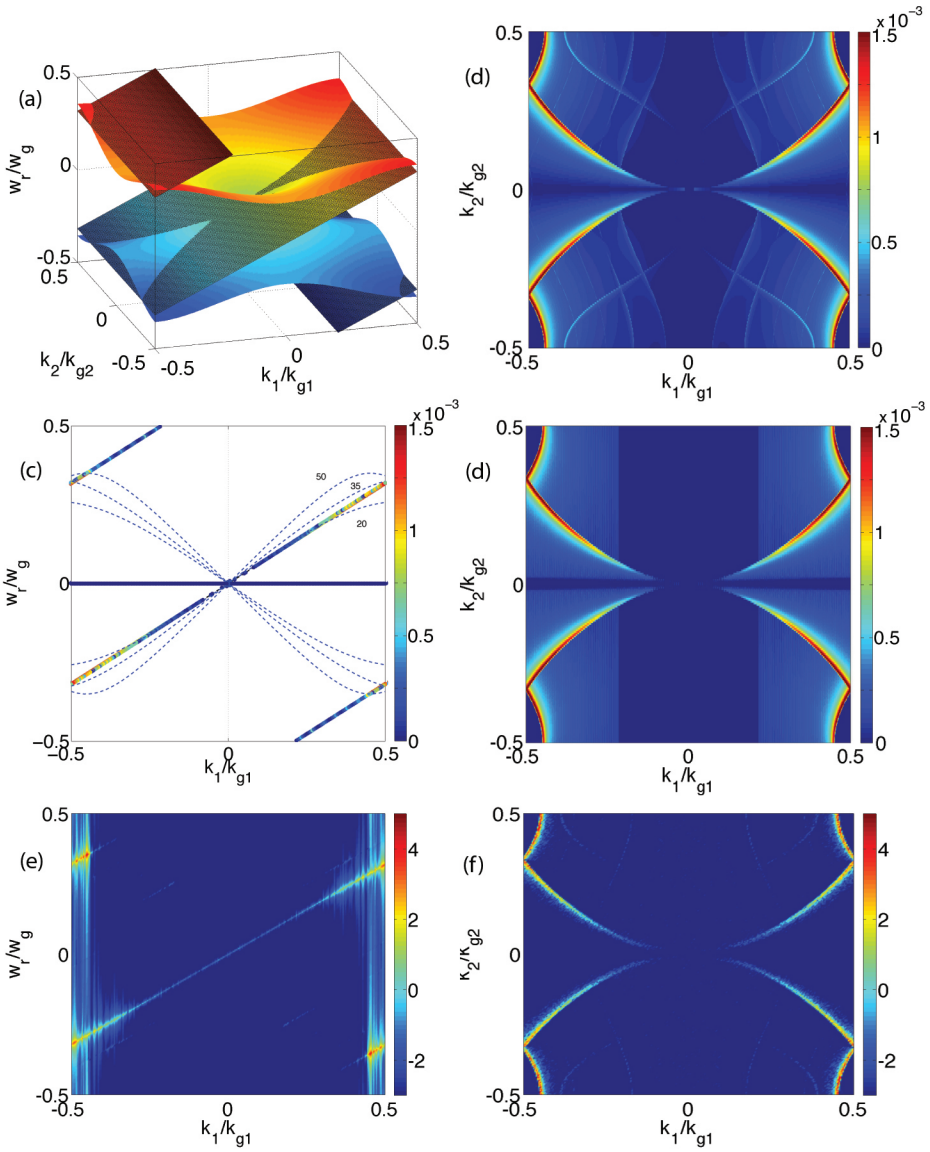

Figure 2.1: Numerical instability pattern in the Yee solver. Growth rates are color-coded, and normalized with ω_g. (a) EM modes intersect with the main beam resonance ($\mu = 0$, $\nu = 0$), and first order space aliasing beam resonances ($\mu = 0$, $\nu_1 = \pm 1$); (b) is the instability pattern ($\mu = 0$, $|\nu_1| \leq 4$) in (k_1, k_2) space; (c) and (d) are the instability pattern ($|\mu| \leq 1$, $|\nu_1| \leq 1$) in (ω_r, k_1) and (k_1, k_2) spaces obtained from solving Eq. (2.21) and (2.22). EM modes for different propagating angles [in degree] and the beam resonances are likewise plotted in (c). (e) presents the corresponding simulation results in (ω_r, k_1) space, and (f) in (k_1, k_2) space.

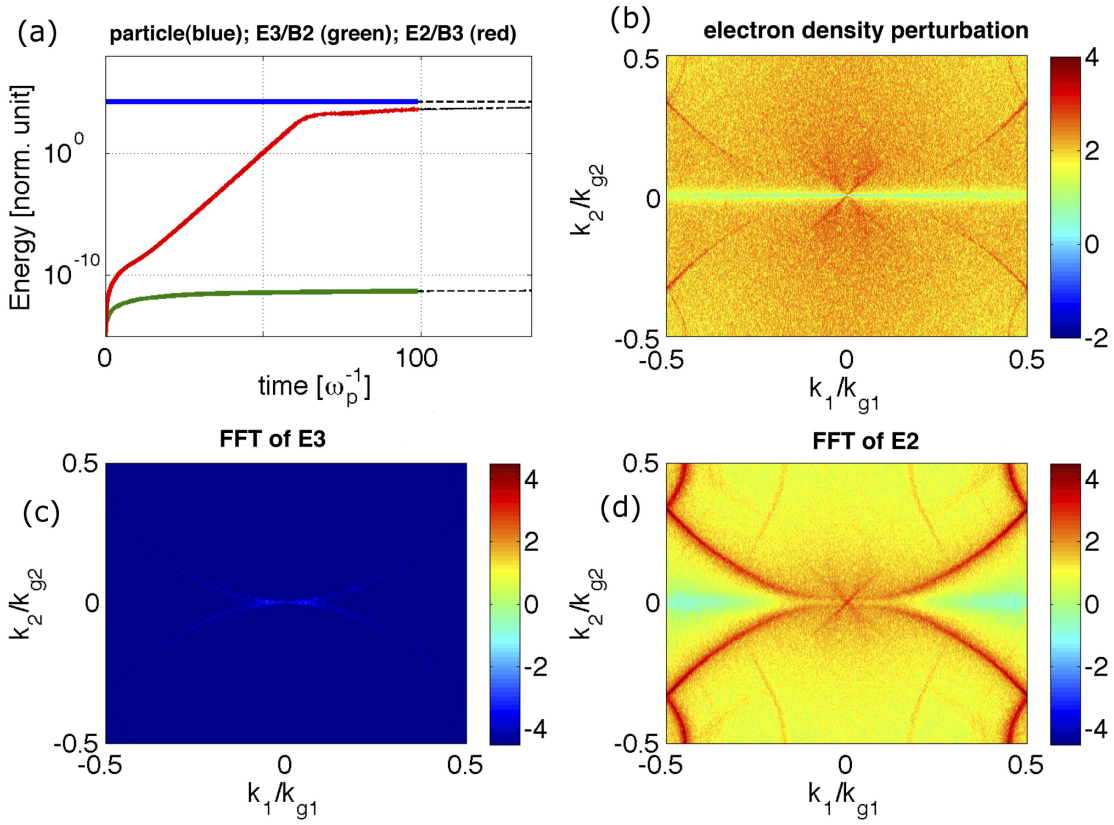

Figure 2.2: We present in (a) the energy evolution of the EM energy for the two cases. The corresponding dotted line indicates their variation in time after $t = 100\ \omega_p^{-1}$; (b) is the plasma electron density perturbation in (k_1, k_2) space. (c) presents the E_3 in (k_1, k_2) space, and (d) presents the E_2 in (k_1, k_2) space.

2.2.2 Simulation comparisons

To compare with the results in the previous section, we conducted simulation studies in the 2D system using the EM-PIC code OSIRIS [30]. In these simulations, a neutral plasma with both the ion and electrons drifting in x_1 at the same relativistic Lorentz factor of $\gamma = 50.0$ is initialized throughout the entire simulation box. Periodic boundary conditions for fields and particles are used. Other parameters for the simulation setup are identical to the numerical solution of dispersion mentioned earlier.

As is shown in figure 2.2 (a), the total EM energy starts to grow violently as the plasma drifts relativistically. The exponential growth indicates that a numerical instability occurs. In addition, the EM field energy in E_2 and B_3 and that in E_3 and B_2 are shown separately. As predicted by the 2D dispersion relation the E_3 and B_2 modes are stable and do not grow. The pattern of E_2 at $t = 100\ \omega_p^{-1}$ is plotted in figure 2.1 (e) and (f), and good agreement for the location and relative amplitude of the unstable modes is obtained when compared against the theoretical prediction [figure 2.1 (c) and (d)].

The EM energy grows with a lower rate after $t = 110\ \omega_p^{-1}$ [figure 2.2 (a)]. The plasma density in this regime is highly modulated by the EM fields. The first order perturbation in plasma electron density [figure 2.2 (b)] shows a similar pattern as for E_2 [figure 2.2 (d)], which confirms they are coupled. Note that no exponential energy growth can be seen in the E_3 field [figure 2.2 (c)]

From the simulation we find that for later times after the instability has evolved into a nonlinear state, the same pattern in (k_1, k_2) space as that of the linear regime still exists. This indicates that the instability will remain near the intersections of the EM modes and beam resonances and that both the linear and nonlinear growth can be mitigated through eliminating or controlling the intersections.

We also carried out a numerical investigation of the 1D dispersion relation Eq. (2.19), and (2.20) using the same simulation parameters as in Tab. 2.2 (with the 1D Courant condition). This confirmed that there is no numerical instability under these conditions which is expected since E_1 is de-coupled from E_2 and E_3 in Eq. (2.19) and each mode is

itself stable.

2.2.3 Parameter scans for minimal instability growth rate

Even before the NCI was analyzed using the numerical dispersion relation, people have found two so-called "magic time steps" under which the NCI growth rate is reduced. This corresponds to $\Delta t = 0.707\Delta x_1$ when using the energy conserving field interpolation [52, 53], and $\Delta t = 0.5\Delta x_1$ when using momentum conserving field interpolation [52]. In the following, we use the dispersion relations in previous sections to confirm these empirical observations.

In figure 2.3, we scanned the grid sizes Δx_1 and time step $\Delta t/\Delta x_1$ for the 2D and 3D Yee solver, and Karkkainen solver [59], and compared the growth rates with the OSIRIS simulations. We have kept $\Delta x_1 = \Delta x_2 (= \Delta x_3)$ during the parameter scan for 2D (and 3D). We likewise presented the OSIRIS simulation data for $\Delta x_1 = 0.1$ together with the asymptotic data. There are several interesting points worth noting in figure 2.3. First, we can see there is an optimized time step [53] $\Delta t_m/\Delta x_1$ where the growth rate is minimized in most cases; on the other hand, the instability growth rate decreases monotonically as the grid sizes increase; second, when the grid sizes are square (2D) or cubic (3D), the optimized time step $\Delta t_m/\Delta x_1$ is an invariant for different Δx_1, in both the momentum conserving (MC) scheme, and energy conserving (EC) scheme; third, the instability growth rate for 2D and 3D are nearly the same for given Δx_1 and $\Delta t/\Delta x_1$ under the same field interpolation scheme; the values for the optimized time steps are also nearly the same in 2D and 3D (note that according to the asymptotic expression, the optimized time step for the Yee solver 3D EC scheme also resides at around $\Delta t_m/\Delta x_1 \approx 0.65$, but we did not plot it out since that Δt_m is beyond the Courant limit for this solver). The parameter scan using the asymptotic expression for the Karkkainen solver with the EC scheme shows the optimized time step at around $\Delta t_m/\Delta x_1 = 0.7$, which agrees with the results reported in Ref. [53]. However, according to our simulation and theoretical results, we found the optimized time step not only in the Karkkainen solver, but also in the Yee solver; and not only for the EC scheme, but also for the MC scheme. This is also reported in Ref. [40] for the 2D cases.

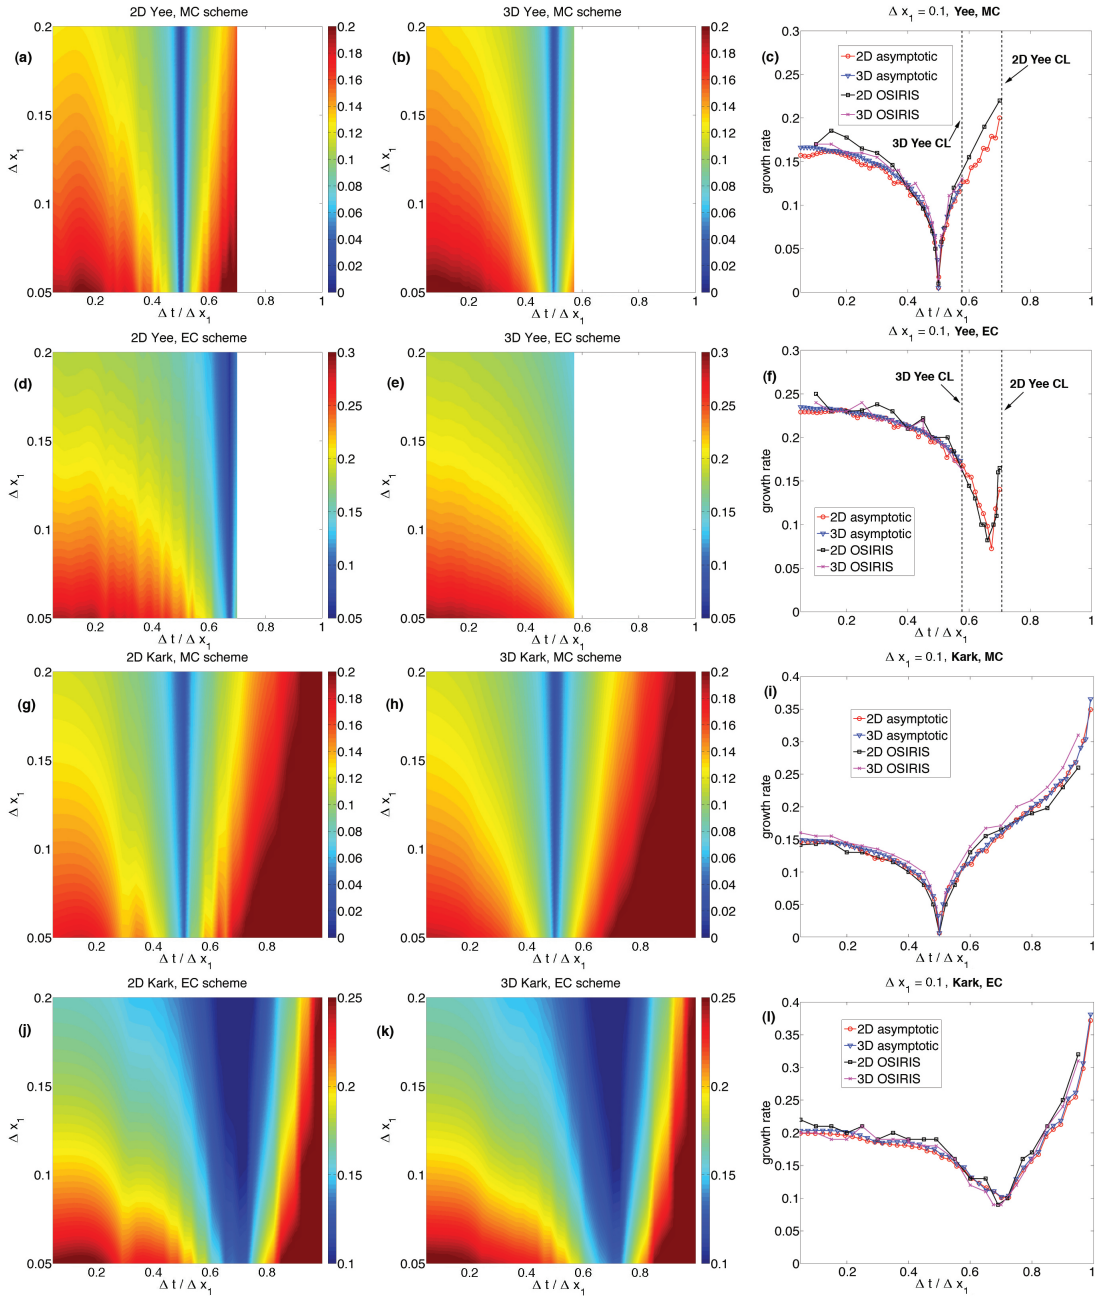

Figure 2.3: Parameter scans of Δx_1 and $\Delta t/\Delta x_1$ for the Yee (first two rows), and Karkkainen (last two rows) solvers. The first and third row uses momentum conserving (MC) scheme, while the second and fourth row uses the energy conserving (EC) scheme. The simulation results are likewise plotted in (c), (f), (i), and (l) at $\Delta x_1 = 0.1$ for comparisons. In (c) and (f) the dotted line at $\Delta t/\Delta x_1 \approx 0.577$ is the 3D Yee solver Courant limit (CL), and that at at $\Delta t/\Delta x_1 \approx 0.707$ is the 2D Yee solver CL.

2.3 Analytical solution of numerical dispersion relations

In this section, we will discuss the analytical approach to systematically obtain the NCI modes from solving Eq. (2.17) and (2.18) [46]. Without losing generality, we will start with its 2D correspondence Eq. (2.21) and (2.22).

As can be seen in Eq. (2.21), when the ω' and k_1' are near the

$$\omega' - k_1' v_0 = (\omega + \mu\omega_g) - (k_1 + \nu_1 k_{g1}) v_0 = 0. \tag{2.23}$$

lines, the numerical solution of Eq. (2.18) for each mode can be analytically obtained by keeping only the corresponding μ and $\vec{\nu}$ terms in Eq. (2.21) since these terms are dominant near the corresponding resonance lines. Note for the cases considered in this disertation, we find it is a good approximation to truncate the sum of ν_2 and only keep the $\nu_2 = 0$ term. For a particular μ and ν_1 the corresponding dispersion relation $\epsilon_{11}\epsilon_{22} - \epsilon_{12}\epsilon_{21} = 0$ for a particular (μ, ν_1) pair becomes

$$\begin{aligned}
&\left([\omega]^2 - [k]_{E2}[k]_{B2} - \frac{\omega_p^2}{\gamma}(-1)^\mu \frac{S_{j1}(S_{E1}[\omega]\omega'/\gamma^2 + S_{B3}v_0^2[k]_{E2}k_2)}{(\omega' - k_1' v_0)^2}\right) \times \\
&\left([\omega]^2 - [k]_{E1}[k]_{B1} - \frac{\omega_p^2}{\gamma}(-1)^\mu \frac{S_{j2}(S_{E2}[\omega] - S_{B3}[k]_{E1}v_0)}{\omega' - k_1' v_0}\right) - \\
&\left([k]_{E1}[k]_{B2} - \frac{\omega_p^2}{\gamma}(-1)^\mu \frac{S_{j1}v_0 k_2(S_{E2}[\omega] - S_{B3}[k]_{E1}v_0)}{(\omega' - k_1' v_0)^2}\right) \times \\
&\left([k]_{E2}[k]_{B1} - \frac{\omega_p^2}{\gamma}(-1)^\mu \frac{S_{j2}S_{B3}v_0[k]_{E2}}{\omega' - k_1' v_0}\right) = 0
\end{aligned} \tag{2.24}$$

After some algebra, Eq. (2.24) can be written as

$$\begin{aligned}
&\left((\omega' - k_1' v_0)^2 - \frac{\omega_p^2}{\gamma^3}(-1)^\mu \frac{S_{j1}S_{E1}\omega'}{[\omega]}\right) \times \\
&\left([\omega]^2 - [k]_{E1}[k]_{B1} - [k]_{E2}[k]_{B2} - \frac{\omega_p^2}{\gamma}(-1)^\mu \frac{S_{j2}(S_{E2}[\omega] - S_{B3}[k]_{E1}v_0)}{\omega' - k_1' v_0}\right) \\
&+ \mathcal{C} = 0
\end{aligned} \tag{2.25}$$

where $\mathcal{C}$ is a coupling term in the dispersion relation

$$\begin{aligned}
\mathcal{C} = \frac{\omega_p^2}{\gamma}\frac{(-1)^\mu}{[\omega]}\Big\{ & S_{j1}S_{E1}\omega'[k]_{E2}[k]_{B2}(v_0^2 - 1) + S_{j2}S_{E2}[k]_{E2}[k]_{B2}(\omega' - k_1' v_0) \\
& + S_{j1}[k]_{E2}(S_{E2}[k]_{B1}k_2 v_0 - S_{B3}k_2 v_0^2[\omega])\Big\}
\end{aligned} \tag{2.26}$$

Much can be learned by investigating Eq. (2.25). First, in the continuous limit ($\Delta t \to 0$, $\Delta x_i \to 0$, and $\mu = \nu_1 = 0$), we have $[\omega] \to \omega$, $S_{E,B} \to 1$, so the coupling term $\mathcal{C}$ vanishes; second, the two factors in the first term of Eq. (2.25) are the Lorentz transformation of the dispersion relation of the Langmuir (longitudinal) mode, and the EM (transverse) mode in a stationary plasma, which in the continuous limit reduce to

$$(\omega - k_1 v_0)^2 - \frac{\omega_p^2}{\gamma^3} = 0 \qquad \omega^2 - k_1^2 - k_2^2 - \frac{\omega_p^2}{\gamma} = 0 \tag{2.27}$$

Consequently, we can identify the numerical Langmuir modes and EM modes for a drifting plasma as

$$(\omega' - k_1' v_0)^2 - \frac{\omega_p^2}{\gamma^3}(-1)^\mu \frac{S_{j1} S_{E1} \omega'}{[\omega]} \approx 0 \tag{2.28}$$

$$[\omega]^2 - [k]_{E1}[k]_{B1} - [k]_{E2}[k]_{B2} - \frac{\omega_p^2}{\gamma}(-1)^\mu \frac{S_{j2}(S_{E2}[\omega] - S_{B3}[k]_{E1} v_0)}{\omega' - k_1' v_0} \approx 0 \tag{2.29}$$

In addition, from Eq. (2.25) we see that when finite grid sizes and time steps are used neither Eq. (2.28) nor Eq. (2.29) leads to instabilty (if the Courant condition is satisfied). Therefore it becomes clear that the NCI is caused by the numerical coupling between modes which are purely longitudinal and purely transverse in the plasma rest frame due to the non-vanishing term $\mathcal{C}$. With the new form Eq. (2.25), we can directly see how the Langmuir mode couples to EM modes. Therefore, reducing or eliminating the coupling term $\mathcal{C}$ is the key to mitigating the NCI. Another interesting fact obtained from Eq. (2.25) is that, if we assume that the ω_p^2 term in Eq. (2.28) and (2.29) are small and can be neglected, when determining the positions of these two modes in Fourier space, the time and space aliasing μ and ν_1 are in the Langmuir modes, while there is no aliasing part in the EM mode. As a side note, it is evident from Eq. (2.26) that in 1D the coupling term vanishes, i.e. $\mathcal{C} = 0$ in the numerical dispersion relation, hence no NCI is found in 1D.

For each pair of (μ, ν_1) there is a corresponding Eq. (2.25). However, in PIC algorithm the range of (ω, k_1) for the quantities defined at discrete locations and time step is limited to the fundamental Brillouin zone $k_i \in (-k_{gi}/2, k_{gi}/2)$, $\omega \in (-\omega_g/2, \omega_g/2)$. As a result, not all the (μ, ν_1) wave-particle resonances exist within the fundamental Brillouin zone. In the following, we describe a way to systematically identify the wave-particle resonance lines

inside the fundamental zone. Taking the parameters in Table 2.2 as an example, we first plot the $(\mu, \nu_1) = (0, 0)$ line [blue line in Fig. 2.4]. As the line extends to the right it meets the boundary of the fundamental zone at $k_1 = 0.5k_{g1}$. To further extend it into the fundamental zone we increase ν_1 by 1, fold the line to the right boundary of k_1, and obtain the $(\mu, \nu_1) = (0, 1)$ line [red line in Fig. 2.4]. The red line extends further until it reaches the $\omega = 0.5\omega_g$ boundary. To extend it further we increase μ by 1, and obtain the $(\mu, \nu_1) = (1, 1)$ line. Additional higher order modes in the fundamental Brillouin zone can be obtained in this way. The negative (μ, ν_1) lines can likewise be obtained by starting from the main Langmuir mode and then extending it to the left, and sets of these (μ, ν_1) lines can be obtained as the lines hit either the boundary at $\omega = -0.5\omega_g$ (μ is reduced by 1), or $k_1 = -0.5k_{g1}$ (ν_1 is reduced by 1). Using the normalization

$$\hat{\omega} + \mu = v_0(\hat{k}_1 + \nu_1)\lambda_1 \tag{2.30}$$

where

$$\hat{\omega} = \frac{\omega}{\omega_g} \qquad \hat{k}_i = \frac{k_i}{k_{gi}} \qquad \lambda_i = \frac{\Delta t}{\Delta x_i} \tag{2.31}$$

the criterion for the Langmuir modes to be inside the fundamental Brillouin zone are $|v_0\lambda_1\nu_1 - \mu| < 0.5 + 0.5v_0\lambda_1$.

Parameters	**Values**
grid size $(k_p\Delta x_1, k_p\Delta x_2)$	$(0.2, 0.2)$
time step $\omega_p\Delta t$	$0.4k_p\Delta x_1$
boundary condition	Periodic
simulation box size (k_pL_1, k_pL_2)	102.4×102.4
plasma drifting Lorentz factor	$\gamma = 50.0$
plasma density	$n/n_p = 100.0$

Table 2.2: Simulation parameters for the 2D relativistic plasma drift simulation. n_p is the reference density, and $\omega_p^2 = 4\pi q^2 n_p/m_e$, $k_p = \omega_p$ (c is normalized to 1).

Note for explicit Maxwell solvers $\lambda_1 < 1$ is a requirement for stable propagation of EM waves in vacuum. The NCI occurs where a resonance line intersects the EM dispersion

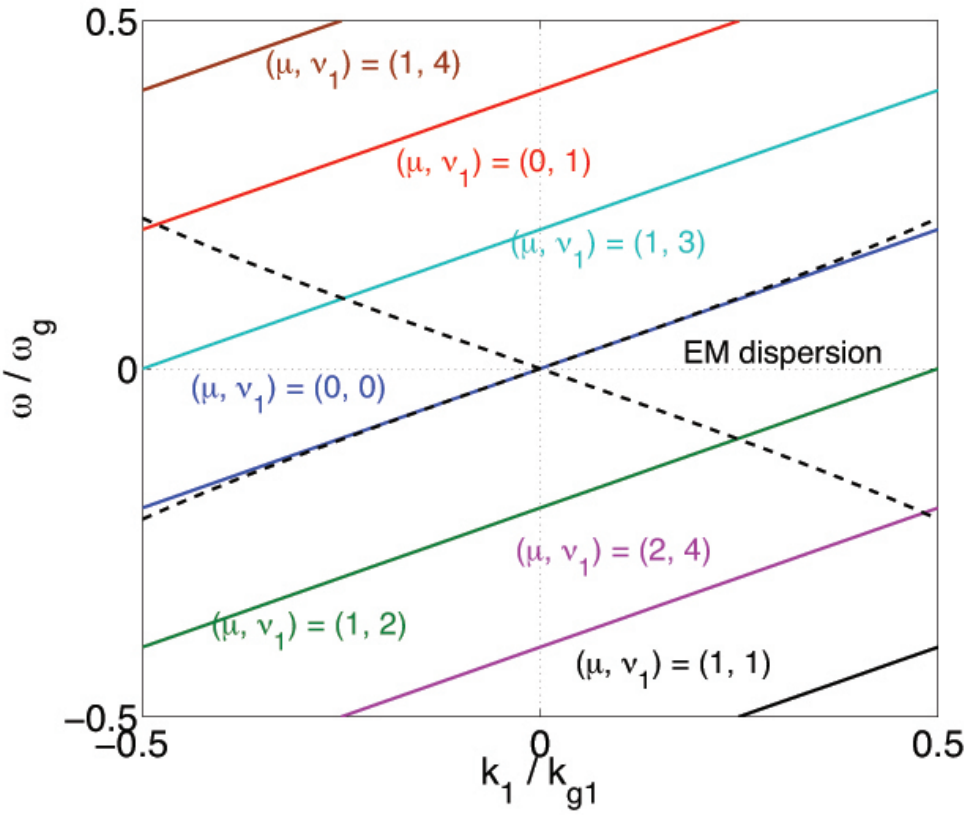

Figure 2.4: The EM dispersion relation together with the beam resonance $\omega' - k_1'\beta = 0$ is shown. The parameters used to plot this figure are listed in Table 2.2.

relation. In Fig. 2.4 we also plot the EM dispersion relation in vacuum as dashed lines. Note for the EM curves we only show $\hat{\omega}$ v.s. $\hat{k}_1$ at $\hat{k}_2 = 0$, but this line varies as $\hat{k}_2$ changes. For the NCI pattern and growth rates associated with each resonance line, we can numerically solve Eq. (2.25) using the corresponding μ and ν_1. Note in [40] a plot similar to Fig. 2.4 can be found (Fig. 1 of Ref. [40]). However, in [40] all the μ are summed over analytically, while in this chapter we emphasize that for a particular resonance line, only one μ term in the elements of $\overleftrightarrow{\epsilon}$ is playing a dominant role. Furthermore, care should be taken when summing over μ and ν_1 as they are not independent sums.

Eq. (2.25) can be used to study the fastest growing mode at $(\mu, \nu_1) = (0, \pm 1)$, yet in here we concentrate on the additional modes. We use Eq. (2.25) to develop analytical expressions within the parameter space we are interested in. Starting from Eq. (2.25), we expand ω' around the beam resonance $\omega' = k_1' v_0$, and write $\omega' = k_1' v_0 + \delta\omega'$, where $\delta\omega'$ is a small term. In addition, we use the relativistic limit $v_0 \to 1$, and expand the finite difference operator $[\omega]$ as

$$[\omega] \approx [\omega]\Big|_{\tilde{k}_1 v_0} + \delta\omega' \frac{\partial[\omega]}{\partial\omega}\Big|_{\tilde{k}_1 v_0} \tag{2.32}$$

where

$$[\omega]\Big|_{\tilde{k}_1 v_0} \equiv \xi_0 = \frac{\sin(\tilde{k}_1 \Delta t/2)}{\Delta t/2} \qquad \frac{\partial[\omega]}{\partial\omega}\Big|_{\tilde{k}_1 v_0} \equiv \xi_1 = \cos(\tilde{k}_1 \Delta t/2) \tag{2.33}$$

where $\tilde{k}_1 = k_1 + \nu_1 k_{g1} - \mu\omega_g$, and $[\omega]^2 \approx \xi_0^2 + 2\xi_0\xi_1\delta\omega'$. In addition, we found it is sufficiently accurate if we neglect the ω^2/γ^3 term in the Langmuir mode in Eq. (2.25). This is why it is essentially the same to say that the instability occurs at wave-particle resonances, beam resonances, or at Langmuir resonances. Moreover, note that ω terms likewise appear in $\vec{S}_B$, and we will separate it from $\vec{S}_B$ by writing

$$\vec{S}_B = \cos(\omega\Delta t/2)\vec{S}_B' \tag{2.34}$$

and expand $\vec{S}_B$ to first order as

$$\vec{S}_B = (\zeta_0 + \zeta_1\delta\omega')\vec{S}_B' \qquad \zeta_0 \equiv \cos(\tilde{k}_1\Delta t/2) \qquad \zeta_1 \equiv -\sin(\tilde{k}_1\Delta t/2)\Delta t/2 \tag{2.35}$$

Using these approximations, we obtain a cubic equation for $\delta\omega'$,

$$A_2\delta\omega'^3 + B_2\delta\omega'^2 + C_2\delta\omega' + D_2 = 0 \tag{2.36}$$

where

$$\begin{aligned}
A_2 =& 2\xi_0^3\xi_1 \\
B_2 =& \xi_0^2\left\{\xi_0^2 - [k]_{E1}[k]_{B1} - [k]_{E2}[k]_{B2} - \frac{\omega_p^2}{\gamma}(-1)^\mu S_{j2}(S_{E2}\xi_1 - \zeta_1 S'_{B3}[k]_{E1})\right\} \\
C_2 =& \frac{\omega_p^2}{\gamma}(-1)^\mu\left\{\xi_0^2 S_{j2}(\zeta_0 S'_{B3}[k]_{E1} - S_{E2}\xi_0) - \xi_1 S_{j1}[k]_{E2}k_2(S_{E2}[k]_{B1} - \zeta_0 S'_{B3}\xi_0)\right. \\
&\left. + \xi_0[k]_{E2}(S_{j2}S_{E2}[k]_{B2} - S_{j1}\zeta_0 S'_{B3}k_2\xi_1 - S_{j1}k_2\zeta_1 S'_{B3}\xi_0)\right\} \\
D_2 =& \frac{\omega_p^2}{\gamma}(-1)^\mu\xi_0[k]_{E2}k_2 S_{j1}\left(S_{E2}[k]_{B1} - \zeta_0 S'_{B3}\xi_0\right)
\end{aligned} \tag{2.37}$$

The coefficients A_2 to D_2 are real, and completely determined by k_1 and k_2. When the discriminant of this cubic equation

$$\Delta = 18A_2B_2C_2D_2 - 4B_2^3D_2 + B_2^2C_2^2 - 4A_2C_2 - 27A_2^2D_2^2 \tag{2.38}$$

satisfies the condition $\Delta < 0$, the cubic equation has one real root and two non-real complex conjugate roots. Therefore, by calculating the discriminant of the cubic equation Eq. (2.38), we can quickly identify the position of the instability for a particular ν_1. We can then use the general formula for the roots of a cubic equation to obtain the growth rate of the corresponding $\vec{k}$ mode. As a result, by solving Eqs. (2.36) and (2.37) we can rapidly calculate the location and growth rate of the instability.

2.4 NCI patterns and growth rates for spectral solver

In section 2.2 we presented the NCI patterns for the Yee solver, by comparing the numerical solution of the dispersion against the simulation results. These patterns can likewise be obtained by using Eqs. (2.36) and (2.37). Nonetheless, to better illustrate the process of analyzing the NCI modes in a systematic way, we will focus on the spectral solver, where the Maxwell's equations are advanced in $\vec{k}$ space rather than $\vec{x}$ space. Specifically, the spatial

derivatives, or correspondingly the spatial finite difference operators $[\vec{k}]$ of the spectral solver are,

$$[k]_i = k_i \tag{2.39}$$

meanwhile, the temporal finite difference operator is the same as for the Yee solver

$$[\omega] = \frac{\sin(\omega\Delta t/2)}{\omega\Delta t/2} \tag{2.40}$$

2.4.1 2D scenario

Without loss of generality, we will start with the 2D case. When using the scheme with spectral solver, we have

$$S_{E1} = S_{E2} = S_{E3} \equiv S_E = S_l \qquad S_{B1} = S_{B2} = S_{B3} \equiv S_B = \cos\frac{\omega\Delta t}{2} S_l.$$
$$S_{j1} = S_{j2} = S_{j3} \equiv S_E = S_l$$

where

$$S_l = \left(\frac{\sin(k_1\Delta x_1/2)}{k_1\Delta x_1/2}\right)^{l+1}\left(\frac{\sin(k_2\Delta x_2/2)}{k_2\Delta x_2/2}\right)^{l+1} \tag{2.41}$$

and l corresponds to the order of the particle shape. Eq. (2.25) reduces to

$$\left((\omega' - k_1' v_0)^2 - \frac{\omega_p^2}{\gamma^3}(-1)^\mu \frac{S_j S_E \omega'}{[\omega]}\right)\left([\omega]^2 - k_1^2 - k_2^2 - \frac{\omega_p^2}{\gamma}(-1)^\mu S_j \frac{S_E[\omega] - S_B k_1 v_0}{\omega' - k_1' v_0}\right)$$
$$+ \frac{\omega_p^2}{\gamma[\omega]}(-1)^\mu S_j k_2^2\{v_0^2(S_E\omega' - S_B[\omega]) - v_0\nu_1 S_E k_{g1}\} = 0 \tag{2.42}$$

And the coefficients A_2, B_2, C_2, and D_2 of Eq. (2.37) become

$$A_2 = 2\xi_0^3\xi_1$$
$$B_2 = \xi_0^2\left\{\xi_0^2 - k_1^2 - k_2^2 - \frac{\omega_p^2}{\gamma}(-1)^\mu S_l^2(\xi_1 - \zeta_1 k_1)\right\}$$
$$C_2 = \frac{\omega_p^2}{\gamma}(-1)^\mu S_l^2\left\{\xi_0^2(\zeta_0 k_1 - \xi_0) + \xi_0 k_2(k_2 - \zeta_0 k_2\xi_1 - k_2\zeta_1\xi_0) - \xi_1 k_2^2(k_1 - \zeta_0\xi_0)\right\}$$
$$D_2 = \frac{\omega_p^2}{\gamma}(-1)^\mu \xi_0 S_l^2 k_2^2(k_1 - \zeta_0\xi_0) \tag{2.43}$$

The coefficients are real and completely determined by k_1 and k_2. We note that this dispersion relation could have been obtained from Eq. (23) if more terms were kept in the expansions for the Q_{ij}.

We now use the cubic equation for $\delta\omega'$ for the coefficients in Eq. (28) to systematically investigate the NCI modes for the spectral solver. In Fig. 2.5 (b), (d), and (f) we present the three sets of modes with the highest growth rate calculated by the analytical expressions Eqs. (2.36) and (2.43), for the parameters listed in Table 2.2, and for linear particle shapes ($l = 1$). Fig. 2.5 (b) shows the unstable $(\mu, \nu_1) = (0, \pm 1)$ modes, which are the fastest growing NCI modes. Fig. 2.5 (d) shows the $(\mu, \nu_1) = (0, 0)$ modes, which have a highly localized pattern of four dots [note that in (d) only one quadrant is plotted]. These modes usually have a maximum growth rate one order of magnitude smaller than the $(\mu, \nu_1) = (0, \pm 1)$ modes. For the parameters listed in Table 2.2, the next fastest growing modes are the $(\mu, \nu_1) = (\pm 1, \pm 2)$ modes which have a maximum growth rate approximately 3 times smaller than the $(\mu, \nu_1) = (0, 0)$ modes (for linear particle shape).

We have performed UPIC-EMMA simulations in 2D to observe various NCI modes in the spectral solver, and to compare with the theory presented above. The simulations use a neutral plasma drifting at a relativistic velocity, with the Lorentz factor $\gamma = 50.0$. The plasma has a uniform initial spatial distribution, and we used the parameters listed in Table 2.2.

In Fig. 2.5 (a), (c), and (e) we show the FFT of the simulation data E_2 at a particular time during the exponential EM energy growth from the NCI. Fig. 2.5 (a) shows results from a simulation with no low-pass filter, and the most prominent modes are those from the $(\mu, \nu_1) = (0, \pm 1)$ resonance. To generate the frames in the middle row, we use a low-pass filter to eliminate the $(\mu, \nu_1) = (0, \pm 1)$ modes. This makes the unstable $(\mu, \nu_1) = (0, 0)$ modes more noticeable. It is shown in Fig. 2.5 (c) that the $(\mu, \nu_1) = (0, 0)$ modes have a highly localized pattern of four dots [in Fig. 2.5 only one quadrant is shown], which agrees with the prediction of the analytic expression. According to Fig. 2.4, there is no intersection between the $(\mu, \nu_1) = (1, 1)$ resonance [or $(\mu, \nu_1) = (-1, -1)$ resonance] and the EM dispersion relation, so the next set of modes of interest are the $(\mu, \nu_1) = (1, 2)$ and

$(\mu, \nu_1) = (-1, -2)$ modes. To make the $(\mu, \nu_1) = (\pm 1, \pm 2)$ modes more noticeable, we use a low-pass filter to filter out the $(\mu, \nu_1) = (0, \pm 1)$ mode, plus a four-dot mask filter to remove the $(\mu, \nu_1) = (0, 0)$ modes. As shown in Figs. 2.5 (e) and (f), the locations of these modes in the simulation agree with the analytic prediction. As a side note, this numerical experiment also shows the simplicity and flexibility of using filters (masks) with complicated shapes in a spectral EM-PIC code to control the unphysical NCI growth.

According to both the theory and simulations, in the parameter space we are interested in, we usually categorize the NCI for a spectral solver into three categories: the fastest growing modes at $(\mu, \nu_1) = (0, \pm 1)$; the second fastest growing modes at $(\mu, \nu_1) = (0, 0)$; and higher order NCI modes with $|\nu_1| > 1$ that have an even smaller growth rate. In the following we will discuss how the locations and positions of these modes change with the simulation parameters.

For the NCI modes with $|\nu_1| \geq 1$, the instability resides around the intersections of the Langmuir mode and EM mode (taking the small time step limit):

$$(1 - v_0^2)k_1^2 + k_2^2 - 2\beta\xi k_1 - \xi^2 = 0 \tag{2.44}$$

where $\xi = \beta\nu_1 k_{g1} - \mu\omega_g$. If we use the normalization in Eq. (2.31) the equations above can be written as (for square cells)

$$(1 - v_0^2)\hat{k}_1^2 + \hat{k}_2^2 - 2\beta\hat{\xi}\hat{k}_1 - \hat{\xi}^2 = 0 \tag{2.45}$$

where $\hat{\xi} = \beta\nu_1 - \mu/\lambda_1$. The positions of the unstable NCI modes in $\vec{\hat{k}}$ space depends only on Δt and Δx_i through their ratio λ_i. Therefore, the position of the $|\nu_1| \geq 1$ NCI does not change if one keeps the ratio of time step to cell size fixed. Moreover, if $\mu = 0$, λ_1 does not appear in Eq. (2.45), which means that the position of the $(\mu, \nu_1) = (0, \pm 1)$ modes are not affected by the time step.

For the NCI at $(\mu, \nu_1) = (0, 0)$, there is no true "intersection" between the corresponding fundamental Langmuir mode and the EM mode [as can be seen by plotting Eq. (2.28) and (2.29) in (ω, k_1) space, see Fig. 2.6 (a)]. However, the two modes interact at highly localized positions determined by the coupling term in Eq. (2.25). To show how the coupling term

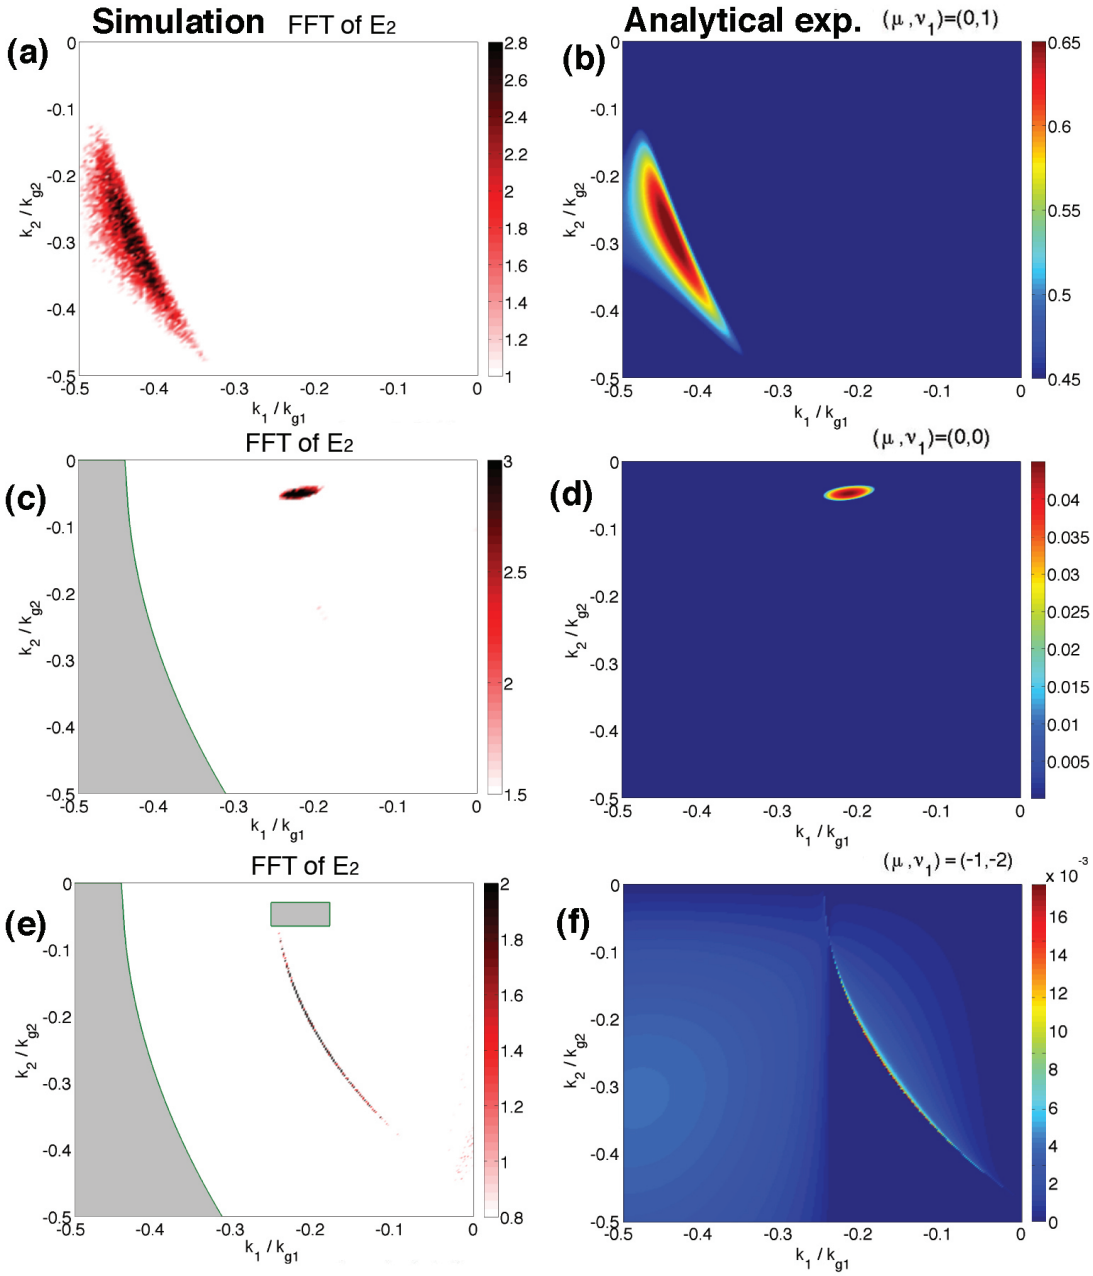

Figure 2.5: (a), (c), and (e) are the FFT of E_2 in the 2D simulations using the parameters listed in Table 2.2. The filter applied in order to observe these modes are illustrated by the grey areas in the plots. (b), (d), and (f) are the corresponding predictions by using the expression Eq. (2.36) and (2.43).

in Eq. (2.25) modifies the Langmuir and EM modes, we plot the solution of Eq. (2.25) at $\hat{k}_1 \approx 0.21$, $-0.07 \leq \hat{k}_2 \leq -0.02$, where the instability is observed. Equation (2.25) is solved both with, and without the coupling term (numerically forcing the coupling term to be zero). The parameters used in solving Eq. (2.25) numerically are the same as in Table 2.2, with $(\mu, \nu_1) = (0, 0)$. It is evident in Fig. 2.6 (a) and (c) that when the coupling term is present, the fundamental Langmuir mode and EM mode are coupled near $-0.057 \leq \hat{k}_2 \leq -0.037$. In Fig. 2.6 (c) where the growth rate is plotted, it becomes clear that in this range of k_2 where the fundamental Langmuir mode and EM mode are coupled, the two modes become complex conjugate pairs with one of them corresponding to instability in this range of $\hat{k}_2$. In Figs. 2.6 (b) and (d), we scan ranges in both $\hat{k}_1$ and $\hat{k}_2$, specifically, we scan the range $\hat{k}_1 \in [-0.28, -0.15]$ and $\hat{k}_2 \in [-0.07, -0.02]$.

We next investigate the sensitivity of the growth rate and location in $\vec{k}$-space to the simulation parameters for the NCI at the fundamental mode $(\mu, \nu_1) = (0, 0)$. Note that we define the position of these modes at the value of $(\hat{k}_1, \hat{k}_2)$ where the growth rate is maximum. In reality there is a range (although highly localized) of modes that go unstable. Fig. 2.7 (a)–(d) shows how the positions and growth rates of those modes change with plasma density and time step. For each simulation setup we plot both the simulation results and the predictions from the analytical expressions. When changing the grid sizes we fix $\Delta x_1 = \Delta x_2$. Fig. 2.7 (a) shows that when the grid sizes increases, the position of the $(\mu, \nu_1) = (0, 0)$ NCI moves farther away from the center of the (k_1, k_2) plot where the interesting real physics resides [red curve in Fig. 2.7 (a)]. We keep Δt constant as Δx_1 changes in Fig. 2.7 (a). The $(\mu, \nu_1) = (0, 0)$ mode also moves farther away from the interesting physics when the time step decreases [see red curve in Fig. 2.7 (b)]. Furthermore, as shown in Fig. 2.7 (c), the growth rate decreases as the time step decreases [blue curve], which is not the case for the fastest growing modes of the NCI. The growth rate also decreases when the grid size increases while keeping Δt fixed [Fig. 2.7 (c) red curve]. When the density of the plasma increases (while fixing $\gamma_b = 50$), the position of the $(\mu, \nu_1) = (0, 0)$ NCI moves away from the center in (k_1, k_2) space, and the growth rates of these modes increase [Fig. 2.7 (d)].

A parameter scan which shows how the growth rate and position of the $(\mu, \nu_1) = (0, 0)$

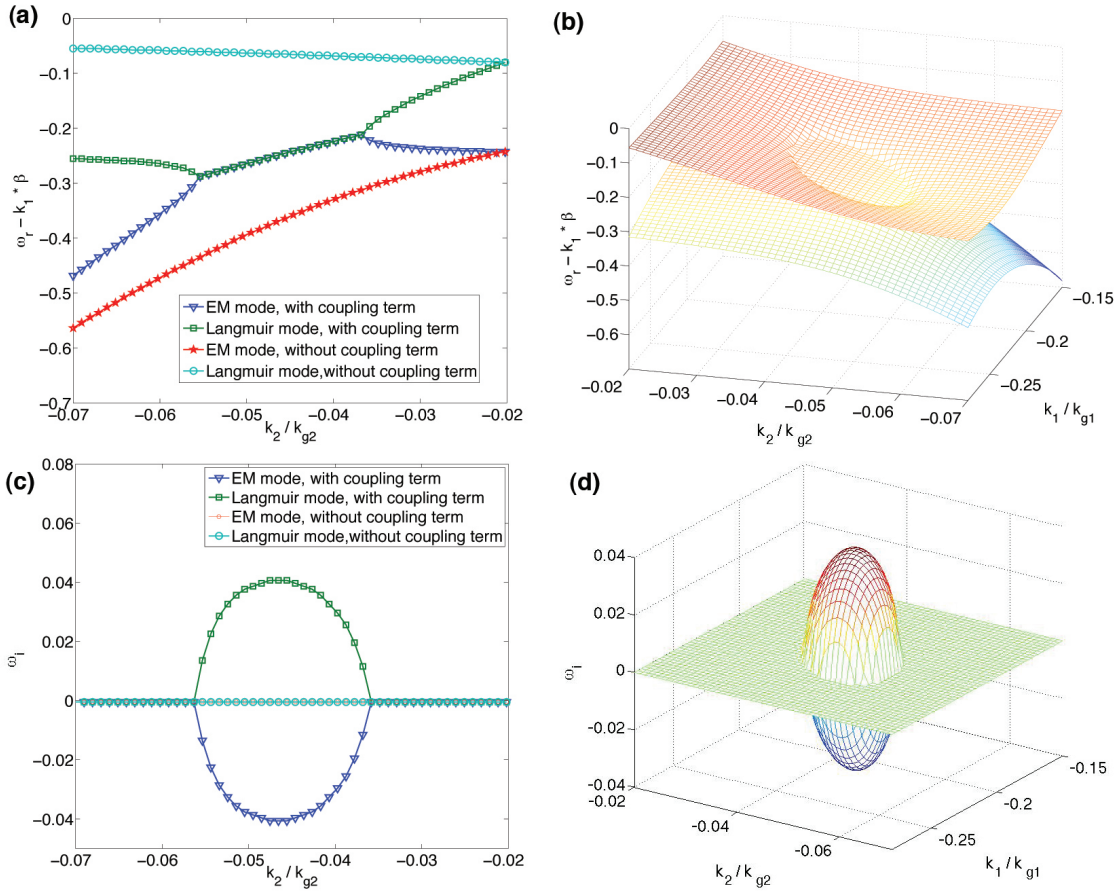

Figure 2.6: Roots of Eq. (2.25) under the parameters listed in Table 2.2. (a) and (c) shows the real, and imaginary parts of the roots between $\hat{k}_1 = 0.21$, and $-0.07 \leq \hat{k}_2 \leq -0.2$, both with and without the coupling terms; meanwhile (b) and (d) shows the real and imaginary part of the roots in the range $-0.28 \leq \hat{k}_1 \leq -0.15$ and $-0.07 \leq \hat{k}_2 \leq -0.02$.

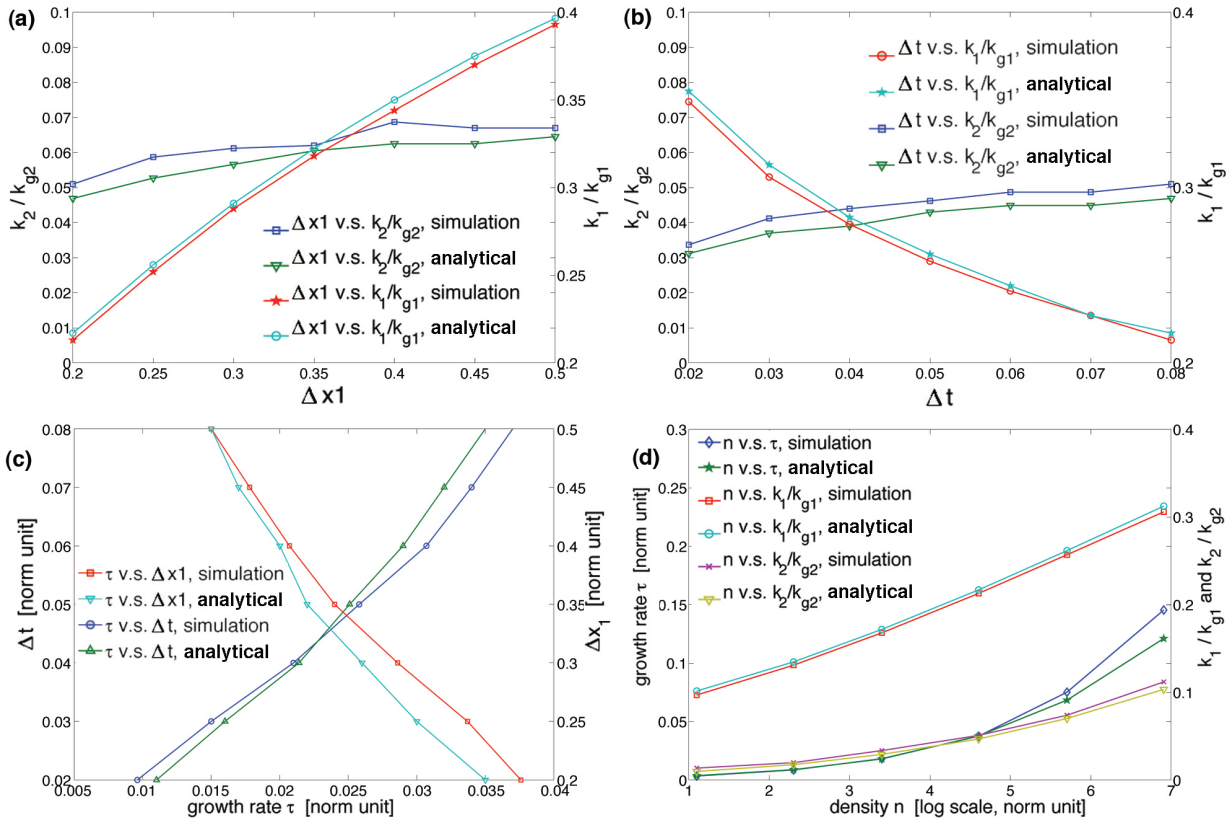

Figure 2.7: Dependence of the position $(\hat{k}_1, \hat{k}_2)$, as well as the growth rate τ of the NCI at the fundamental Langmuir mode to grid sizes Δx_1 (with $\Delta x_1 = \Delta x_2$ fixed), time step Δt, and plasma density n_p.

modes change with different choices of the grid sizes and time step, is shown in Fig. 2.8. Note that we are keeping $\Delta x_1 = \Delta x_2$ in the parameter scan. By examining Fig. 2.8, we see that by reducing the $\Delta t/\Delta x_1$ ratio, the instability at the fundamental Langmuir mode moves towards larger $\hat{k}_1$ and the growth rate decreases. This is a unique characteristic of the $(\mu, \nu_1) = (0, 0)$ modes, i.e., the growth rate of the fastest growing modes does not decrease as $\Delta t/\Delta x_1$ decreases. This is illustrated in Fig. 2.9 (c) where the growth rate of the $(\mu, \nu_1) = (0, 0)$ and $(0, 1)$ modes are plotted against $\Delta t/\Delta x_1$ for $k_p \Delta x_1 = 0.2$. When the fastest growing modes are filtered out in a simulation, if the grid size is restricted to resolve the characteristic length of physical modes, the position of the $(\mu, \nu_1) = (0, 0)$ mode can be moved to larger $\hat{k}_1$ by simply using a smaller time step.

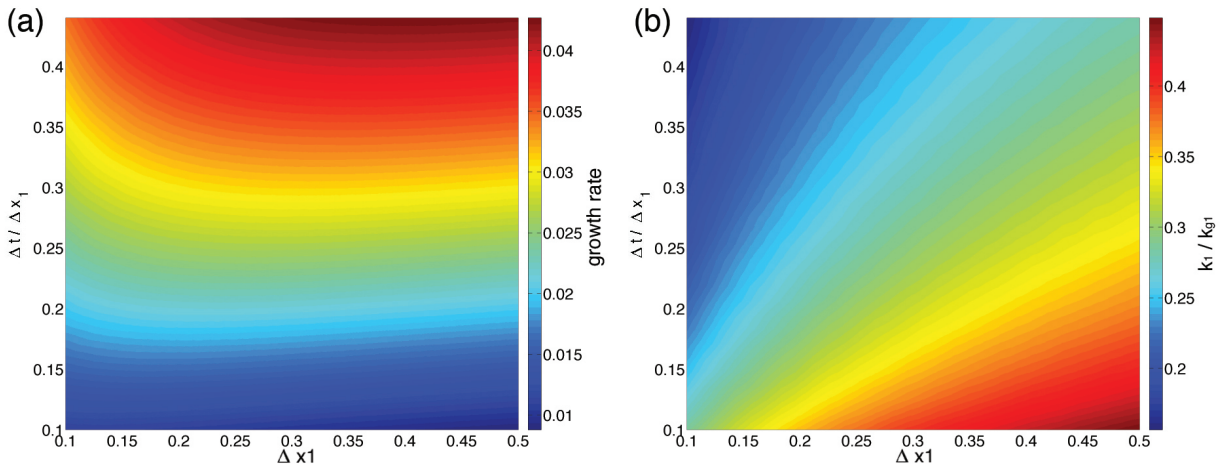

Figure 2.8: Dependence of the (a) growth rate, and (b) k_1 position of NCI at the fundamental Langmuir mode for grid sizes $0.1 \leq \Delta x_1 \leq 0.5$ (with $\Delta x_1 = \Delta x_2$ fixed), and $0.1 \leq \Delta t/\Delta x_1 \leq 0.45$.

Meanwhile, when the time step is fixed, the growth rates of higher order NCI ($|\nu_1| > 1$) unstable modes can be efficiently reduced by using higher order particle shapes. In Fig. 2.9 (a) we show how using different particle shapes changes the growth rate of the various NCI modes. The parameters in Table 2.2 are used for this figure. The result indicates that, while using higher order particle shapes is very efficient in reducing the growth rate of higher order NCI modes, it is less efficient for the $(\mu, \nu_1) = (0, 0)$ mode. We also compared

results with different grid sizes (while fixing $\Delta t/\Delta x_1 = 0.4$), as shown in Fig. 2.9 (b). It indicates that reducing the grid size (while fixing $\Delta t/\Delta x_1$) helps reduce the growth rate of the $(\mu, \nu_1) = (0, 0)$ mode, but not for the modes with $\nu_1 \neq 0$.

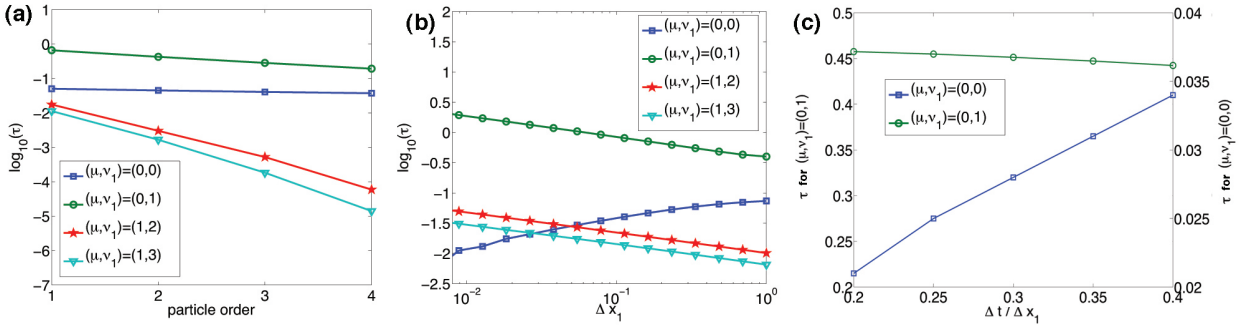

Figure 2.9: (a) shows the dependence of the growth rate on particle shapes. (b) shows the dependence of the growth rate on grid size for various ν_1 modes. (c) shows the dependence of the growth rate on time step when the grid sizes are fixed. Parameters listed in Table 2.2 are used for these plots.

2.4.2 3D scenario

We next discuss the NCI in three dimensions for a spectral solver. Based on the results for the 2D case, we write the full dispersion relation into the coupling between a Langmuir and an EM mode. For the spectral solver, the dispersion relation for a specific μ, ν_1 mode can be rewritten as,

$$\left([\omega]^2 - k_1^2 - k_2^2 - k_3^2 - \frac{\omega_p^2}{\gamma}(-1)^\mu S_j \frac{S_E[\omega] - S_B k_1 v_0}{\omega' - k_1' v_0}\right)$$
$$\left\{\left((\omega' - k_1' v_0)^2 - \frac{\omega_p^2}{\gamma^3}(-1)^\mu \frac{S_j S_E \omega'}{[\omega]}\right)\left([\omega]^2 - k_1^2 - k_2^2 - k_3^2 - \frac{\omega_p^2}{\gamma}(-1)^\mu S_j \frac{S_E[\omega] - S_B k_1 v_0}{\omega' - k_1' v_0}\right)\right.$$
$$\left. + \frac{\omega_p^2}{\gamma[\omega]}(-1)^\mu S_j (k_2^2 + k_3^2)\{v_0^2(S_E\omega' - S_B[\omega]) - v_0 \nu_1 S_E k_{g1}\}\right\} = 0 \tag{2.46}$$

For the instability mode near the resonance line, we can assume

$$\left([\omega]^2 - k_1^2 - k_2^2 - k_3^2 - \frac{\omega_p^2}{\gamma}(-1)^\mu S_j \frac{S_E[\omega] - S_B k_1 v_0}{\omega' - k_1' v_0}\right) \neq 0 \tag{2.47}$$

$$\left((\omega' - k_1' v_0)^2 - \frac{\omega_p^2}{\gamma^3}(-1)^\mu \frac{S_j S_E \omega'}{[\omega]}\right)\left([\omega]^2 - k_1^2 - k_2^2 - k_3^2 - \frac{\omega_p^2}{\gamma}(-1)^\mu S_j \frac{S_E[\omega] - S_B k_1 v_0}{\omega' - k_1' v_0}\right)$$
$$+ \frac{\omega_p^2}{\gamma[\omega]}(-1)^\mu S_j (k_2^2 + k_3^2)\{v_0^2(S_E\omega' - S_B[\omega]) - v_0 \nu_1 S_E k_{g1}\} = 0 \tag{2.48}$$

which, as in the 2D case, can be viewed as the coupling between the Langmuir and EM mode. When Eq. (2.48) is compared with Eq. (2.42), we can see that for the spectral solver the equations in 3D can be obtained by replacing k_2^2 with $k_2^2 + k_3^2$ in its 2D counterpart. As a result, the pattern of instability in 3D can be conveniently deduced. The location of the $(\mu, \nu_1) = (0, 0)$ NCI modes in $\vec{k}$ space in 3D can be obtained as follows. Pick a point in (k_2, k_3) space, then the growth rate and location in k_1 space of this mode will be the same as for $k_2^{2D} = \sqrt{(k_2^{3D})^2 + (k_3^{3D})^2}$, where $k_{1,2,3}^{3D}$ and $k_{1,2}^{2D}$ are the coordinates of the modes in the 3D and 2D scenario respectively (assuming $\Delta x_1^{3D} = \Delta x_2^{3D} = \Delta x_3^{3D} = \Delta x_1^{2D} = \Delta x_2^{2D}$ and $\Delta t^{3D} = \Delta t^{2D}$). This indicates that the unstable modes form a "ring" pattern in the (k_2, k_3) space at specific values of $\hat{k}_1$. In addition, the maximum growth rate τ of these modes has $\tau^{2D} \approx \tau^{3D}$ when $\Delta x_1^{3D} = \Delta x_2^{3D} = \Delta x_3^{3D} = \Delta x_1^{2D} = \Delta x_2^{2D}$ and $\Delta t^{3D} = \Delta t^{2D}$.

In Fig. 2.10 we present data from a 3D simulation of a drifting plasma, using the same parameters as in Table 2.2 except now $\Delta t^{3D} = 0.35\Delta x_1$ (so that the Courant condition is satisfied). The different values of Δt and the different noise sources in 3D v.s. 2D means the results will not be identical. We plot the FFT of E_2 in each panel. In Fig. 2.10 (a) the real frequency v.s. $\hat{k}_1$ are plotted along with the line $\hat{\omega}_r = \hat{k}_1 \beta$. This data was obtained for a line out along x_1 located at the middle of the box. In Fig. 2.10 (b) a 3D plot of E_2 in $\vec{k}$-space is shown at a time during the exponential growth (before saturation). Only modes with amplitudes above 1/30 of the maximum mode are plotted. The predicted rings are clearly present. In Fig. 2.10 (c) and (d) cross sections of the plot in Fig. 2.10 (b) are shown.

In analogy with the 2D case, we can filter out this instability by applying a mask filter to eliminate the corresponding modes in the ring. For example we have a mask that blocks

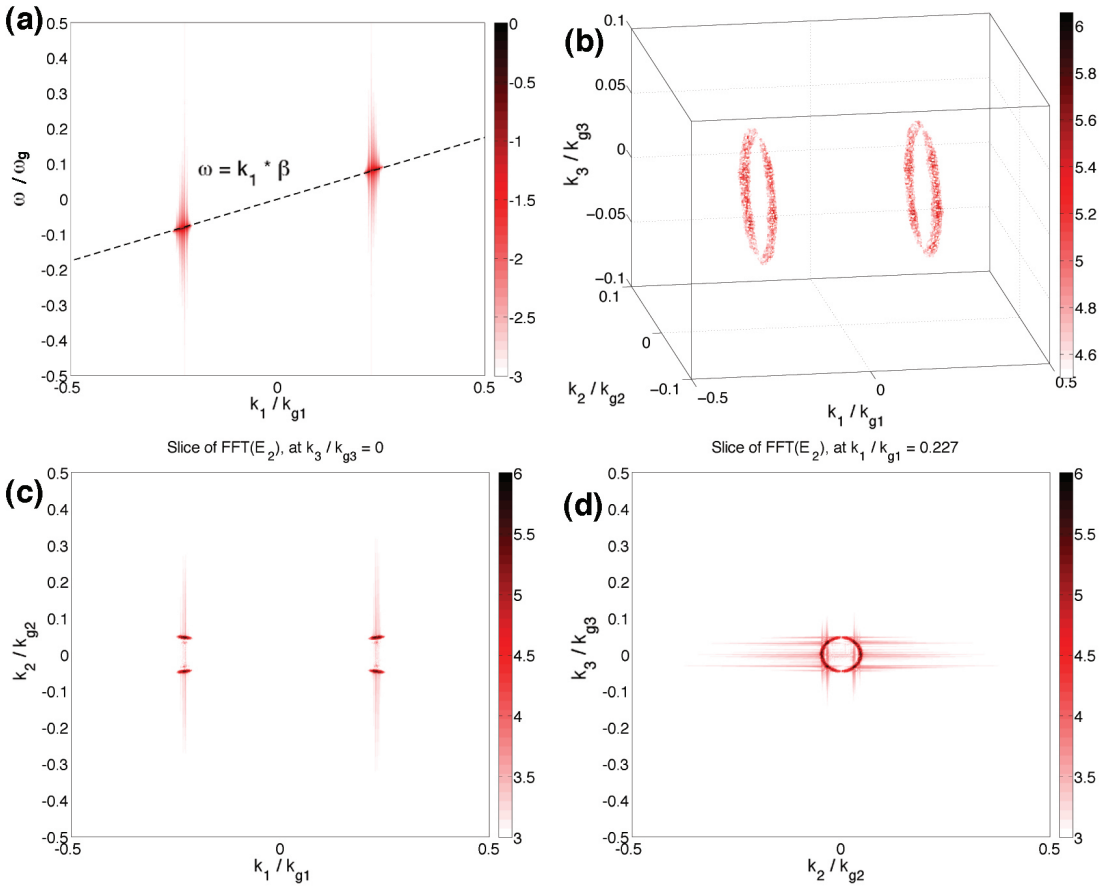

Figure 2.10: This figure shows the dominant NCI modes after the fastest growing modes are filtered out for a 3D simulation. (a) shows in 3D the $(\mu, \nu_1) = (0, 0)$ mode which resides at the main resonance $\omega = k_1\beta$; (b), (c), and (d) are the positions of this NCI mode in $\vec{\hat{k}}$ space.

out all the modes between

$$0.175 \leq \hat{k}_1 \leq 0.275 \qquad 0.027^2 \leq \hat{k}_2^2 + \hat{k}_3^2 \leq 0.067^2$$

A parameter scan of the growth rate and position of the unstable modes using UPIC-EMMA, as well as comparison between the analytical predictions is presented in Fig. 2.11 (a)–(d). The variable k_R in these plots refers to the radius of the ring pattern. It confirms that when comparing the unstable mode for a 2D case (see Fig. 2.7) against its counterpart in 3D, we have approximately $\hat{k}_1^{3D} = \hat{k}_1^{2D}$, and $\hat{k}_R^2 \equiv (\hat{k}_2^{3D})^2 + (\hat{k}_3^{3D})^2 = (\hat{k}_2^{2D})^2$ (where $\hat{k}_R = k_R/k_{g1}$), and $\tau^{3D} \approx \tau^{2D}$. Therefore, the 3D NCI can effectively be eliminated by using the same strategies as in 2D.

2.5 Summary

In this section we start by deriving the general numerical dispersion relation of plasma in a PIC system, and then obtain the detailed expressions for the dispersion tensor for a cold relativistically drifting plasma. The corresponding NCI modes can be directly solved from the dispersion, which agrees well with our simulation observations. Alternatively, the NCI modes can be analytically resolved mode by mode. The idea is that for a specific spatial and temporal aliasing, only one term will be dominant. This significantly simplifies the expression of the dispersion tensor, which allows us to write out the dispersion relations in a compact form Eq. (2.25).

The significance of Eq. (2.25) is two-folded. Firstly, it gives a direct physical interpretation of the source of NCI: the NCI is the unphysical coupling between the EM modes and main and aliasing Langmuir modes due to the use of finite spatial grids and time step. Secondly, it provides an accurate way of deriving the analytical solutions for rapid evaluations of the pattern and growth rates of the NCI modes Eqs. (2.36) and (2.37). As we can see in later chapters, these two equations become the workhouse when we are modifying the Maxwell solver to eliminate the NCI modes.

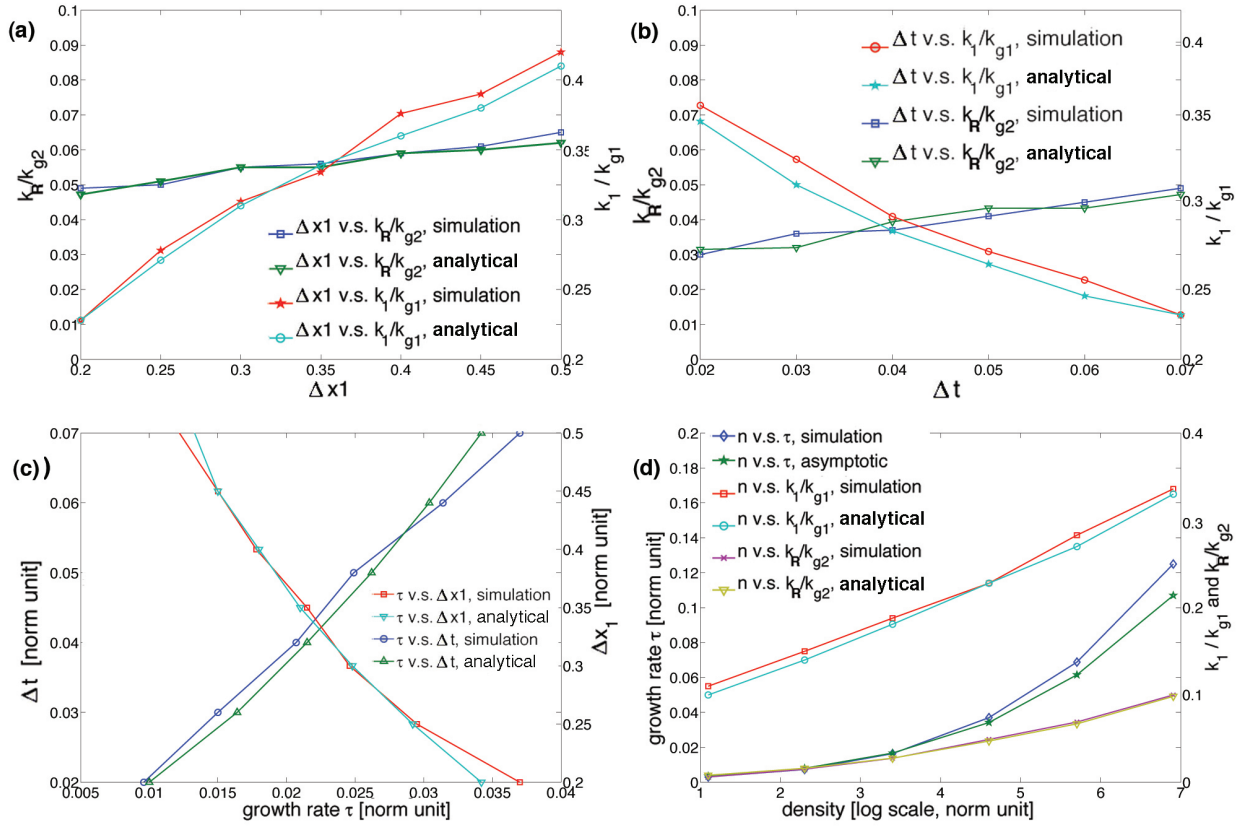

Figure 2.11: Dependence of the position $(\hat{k}_1, \hat{k}_R)$, as well as the growth rate τ of the NCI at the fundamental Langmuir mode to grid sizes Δx_1 (with $\Delta x_1 = \Delta x_2$ fixed), time step Δt, and plasma density n_0 in 3D. The variable k_R refers to the radius of the ring pattern of $(\mu, \nu_1) = (0, 0)$ NCI modes in 3D.

CHAPTER 3

Elimination of NCI in multi-dimensional spectral solver

The theory developed in Chapter 2 provides us with solid evidence that the NCI is due to the unphysical coupling between EM modes and Langmuir modes (both main and aliasing). Writing down the Langmuir mode in the approximate form as

$$\omega' - k_1' v_0 \approx 0 \tag{3.1}$$

we can see that for whatever PIC algorithm we apply, the form of the Langmuir modes will stay the same. It is clear that the Langmuir mode is indistinguishable from wave particle (or beam) resonances to the lowest order. On the other hand, the EM dispersion

$$[\omega]^2 - [k]_{E1}[k]_{B1} - [k]_{E2}[k]_{B2} - \frac{\omega_p^2}{\gamma}(-1)^{\mu}\frac{S_{j2}(S_{E2}[\omega] - S_{B3}[k]_{E1}v_0)}{\omega' - k_1' v_0} \approx 0 \tag{3.2}$$

are flexible in the sense that $[\vec{k}]_{E,B}$ can be easily modified when an alternative spatial finite difference form is applied when solving the Maxwell equation. In Chapter 2 we have already shown how the patterns of the NCI varies with the EM dispersion relation, i.e. it varies when we switch from a Yee solver to a spectral solver. It is then reasonable to argue that the design of a NCI-friendly Maxwell solver is the key to develop a NCI elimination scheme. In fact Chapter 2 already showed that when compared with the Yee solver, the spectral solver shows great potential in eliminating the NCI as the fastest growing NCI modes are on the edges of the fundamental Brillouin zone, while the growth rates for the main NCI modes are already one order of magnitude smaller than the fastest growing modes. We note that focusing on eliminating the coupling term $\mathcal{C}$ can potentially be another path for eliminating the NCI.

In this chapter, we discuss the NCI elimination scheme for a spectral solver that advances the Maxwell equations in $\vec{k}$ space. The validity of this scheme proves to be extremely effective

by the clean simulation results from our production spectral PIC code UPIC-EMMA which implements this elimination scheme. Note that in the following chapters, we will present our design for new Maxwell solvers for NCI elimination. The elimination methods for all these new Maxwell solvers are based on the elimination scheme described in this chapter.

3.1 Strategies for eliminating NCI

Based on the behavior of the unstable NCI modes, we now discuss approaches for controlling it. Once the NCI is adequately controlled, high fidelity simulations of relativistically drifting plasma can be carried out. The new form for the dispersion relation in Eq. (2.25) can also be used to investigate the NCI for a variety of solvers.

The approach we use to eliminate the NCI is to first move the unstable modes to large $\vec{\hat{k}}$'s that are outside the region in $\vec{\hat{k}}$ space where important physics is occurring. As discussed in Chapter 2, for the spectral solver the fastest growing modes at $(\mu, \nu_1) = (0, \pm 1)$ exist at large $|\vec{\hat{k}}|$ (the edge of the fundamental Brillouin zone). In addition, as discussed earlier their location in $\vec{\hat{k}}$-space does not change much as the grid sizes (for square or cubic cells) and time step are varied.

As shown in Fig. 2.5, the second fastest growing mode at $(\mu, \nu_1) = (0, 0)$ is highly localized in $\vec{\hat{k}}$-space and can be removed through a mask filter. However, these modes may exist near modes of physical interest, and for LWFA boosted frame simulations the plasma only exists in a small region of the simulation window. For such situations simply applying a mask filter may also effect the physics. We therefore eliminate those modes by first reducing the time step (while keeping the cell size fixed). As shown in Figs. 2.8 (a) and (b), this both moves the unstable modes to higher $\hat{k}_1$ and lowers the growth rate.

To investigate how reducing the time step changes the NCI, 2D simulations using the same parameters as those shown in Fig. 2.5, but with a reduced time step of $\Delta t = 0.1\Delta x_1$ are conducted. The corresponding intersection between beam resonances and the EM wave for this time step are illustrated in Fig. 3.1 (a), while the corresponding simulation data and analytical prediction for $\Delta t = 0.1\Delta x_1$ are shown in Fig. 3.1 (c)–(f). From Fig. 3.1 (c)

and (e) we see as expected that when the time step is reduced, the growth rate and pattern of the fastest growing modes at $(\mu, \nu_1) = (0, \pm 1)$ do not change much [compared with Fig. 2.5 (a) and (b)]. However, for the $(\mu, \nu_1) = (0, 0)$ modes shown in Fig. 3.1 (d) and (f) , the locations move away from the center [compared with Fig. 2.5 (c) and (d)], while the growth rate is reduced by approximately a factor of 4.

In addition, when the time step is reduced to suppress the $(\mu, \nu_1) = (0, 0)$ mode, the locations and growth rate for the higher order $|\nu_1| > 1$ modes also change. As seen in Fig. 3.1 (a), the next aliasing beam resonance after $(\mu, \nu_1) = (0, \pm 1)$ is $(\mu, \nu_1) = (0, \pm 2)$ rather than $(\mu, \nu_1) = (1, \pm 1)$. It is easy to see that in this case the $(\mu, \nu_1) = (0, \pm 2)$ resonance line has no intersection with the EM mode in the fundamental Brillouin zone. The beam resonance line for an intermediate time step of $\Delta t = 0.225 \Delta x_1$ is likewise shown in Fig. 3.1 (b). This illustrates how gradually reducing the time step changes the NCI modes in the fundamental zone.

Reducing the time step in FFT-based Maxwell solvers is preferable in relativistically drifting plasma simulation as it not only provides better NCI properties, but also provides better accuracy to the Maxwell solver, and pusher in the algorithm. However, it comes at a cost of increased computational loads. In the following we describe another approach for eliminating the $(\mu, \nu_1) = (0, 0)$ NCI modes for the spectral solver with a minor modification in the EM dispersion curve. This approach can be used alone, or combined with the reduced time step to achieve complete elimination of the $(\mu, \nu_1) = (0, 0)$ NCI modes.

As seen from Fig. 2.6, the $(\mu, \nu_1) = (0, 0)$ NCI modes = to the intersection between EM mode and the main Langmuir modes, but rather a coupling when the two curves get near each other at localized region in $\vec{k}$ space. To eliminate this coupling region, we now artificially create a small bump to the EM mode by slightly modifying the corresponding $[k]_1$ operator in the Maxwell solver

$$[k]_1 = k_1 + \Delta k_{mod} \tag{3.3}$$

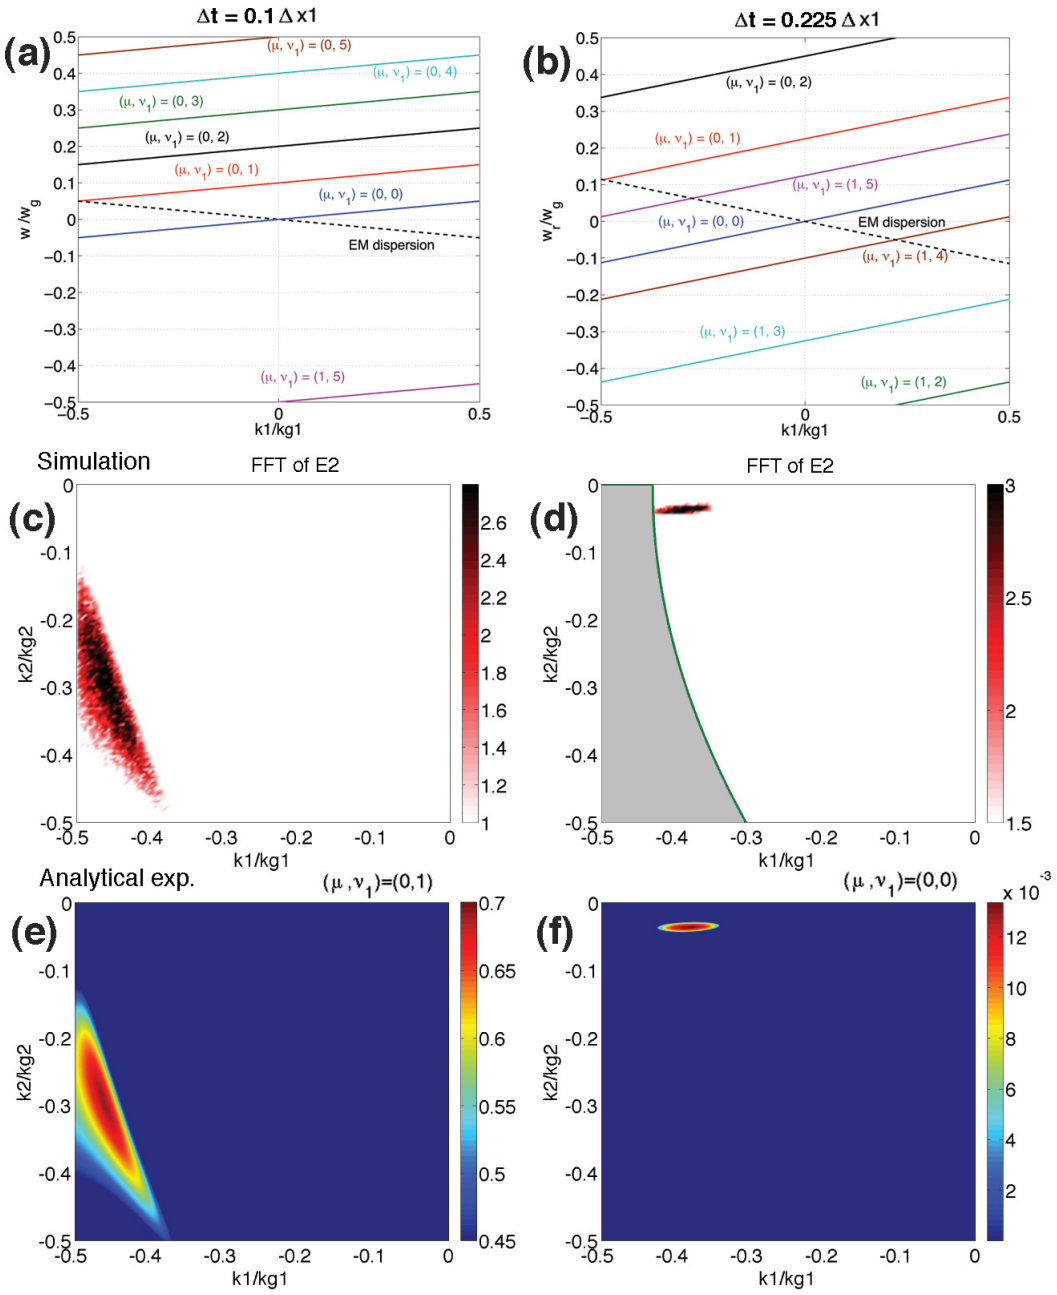

Figure 3.1: (a) and (b) show the EM dispersion relation together with the beam resonance $\omega' - k_1'\beta = 0$, for $\Delta t = 0.1\Delta x_1$ and $\Delta t = 0.225\Delta x_1$ (and other parameters the same as listed in Table 2.2). (c) and (d) are the FFT of E_2 in the corresponding 2D simulations. The filter applied in order to observe the $(\mu, \nu_1) = (0, 0)$ mode is illustrated by the grey areas in (d). (e) and (f) are the corresponding analytical predictions by using the expression Eq. (2.36) and (2.43).

where

$$\Delta k_{mod} = \Delta k_{mod,\max} \cos^2\left(\frac{k_1 - k_{1m}}{k_{1,\min} - k_{1,\max}}\frac{\pi}{2}\right)\cos^2\left(\frac{k_2}{k_{2,\max}}\frac{\pi}{2}\right) \tag{3.4}$$

in the range $k_{1,\min} < |k_1| < k_{1,\max}$, and $\Delta k_{mod} = 0$ otherwise. $k_{1,\min}$, $k_{1,\max}$, and $k_{inc,\max}$ are determined by the $(\mu, \nu_1) = (0, 0)$ NCI modes to be eliminated, and $k_{1m} = (k_{1,\min} + k_{1,\max})/2$.

Consider the drifting plasma simulation with $\Delta t = 0.4\Delta x_1$ and other simulation parameters are listed in Table 2.2 as an example. In Fig. 3.2 (a), (c), and (d) we illustrate how the EM dispersion (in vacuum) would change as we apply this modification to the $[k]_1$ operator in the solver in order to eliminate the $(\mu, \nu_1) = (0, 0)$ NCI modes completely. In Fig. 3.2 (c) we show the distribution of $|\omega - \bar{\omega}|/\omega_g$, to indicate how the EM dispersion is modified in the fundamental Brillouin zone, where ω and $\bar{\omega}$ are the frequency corresponding to a particular (k_1, k_2) in the original, and revised EM dispersion, respectively. In Fig. 3.2 (a) and (d) we show the corresponding EM dispersion for $(\omega, \hat{k}_2)$ and $(\bar{\omega}, \hat{k}_2)$ at $\hat{k}_1 = 0.205$, and $(\omega, \hat{k}_1)$ and $(\bar{\omega}, \hat{k}_1)$ at $\hat{k}_2 = 0$ respectively (as the lines of $\hat{k}_1 = 0.205$ and $\hat{k}_2 = 0$ cross the point where the maximum value of Δk_{mod} is reached) to show how much the dispersion is modified. When substituting this $[k]_1$ operator in Eqs. (2.36) and (2.37) while keeping $[k]_2 = k_2$, we can see there are is unstable root for $(\mu, \nu_1) = (0, 0)$, i.e. when the modified $[k]_1$ operator is used in the solver, there is no $(\mu, \nu_1) = (0, 0)$ NCI mode predicted by the theory. In this case $k_{1,\min}/k_{g1} = 0.15$, $k_{1,\max}/k_{g1} = 0.26$, $k_{2,\max}/k_{g2} = 0.125$, and $k_{inc,\max}/k_{g1} = 0.0095$. In Fig. 3.2 (b) we plot the growth in energy for E_2 for the cases with $\Delta t = 0.4\Delta x_1$ and $\Delta t = 0.2\Delta x_1$, as well as the case with $\Delta t = 0.4\Delta x_1$ plus the EM dispersion relation modification. In all these cases a low-pass filter is used to eliminate the fastest growing $(\mu, \nu_1) = (0, \pm 1)$ modes. As shown in Fig. 3.2 (b) for the blue ($\Delta t = 0.4\Delta x_1$), and red ($\Delta t = 0.2\Delta x_1$) curve, the exponential energy growth is due to the $(\mu, \nu_1) = (0, 0)$ modes; meanwhile in the case where the EM dispersion modification is applied (black curve), the energy growth due to $(\mu, \nu_1) = (0, 0)$ modes is completely eliminated. Note later in time the energy grows exponentially (with a much lower growth rate, not shown in the plot) due to the higher order modes $(\mu, \nu_1) = (\pm 1, \pm 2)$. In these simulations we used second order particle shape. As discussed earlier in section 2, if one needs to further suppress the NCI by

reducing the growth rate of the $(\mu, \nu_1) = (\pm 1, \pm 2)$ NCI modes, one can use a higher order particle shape as discussed in section 2.4.

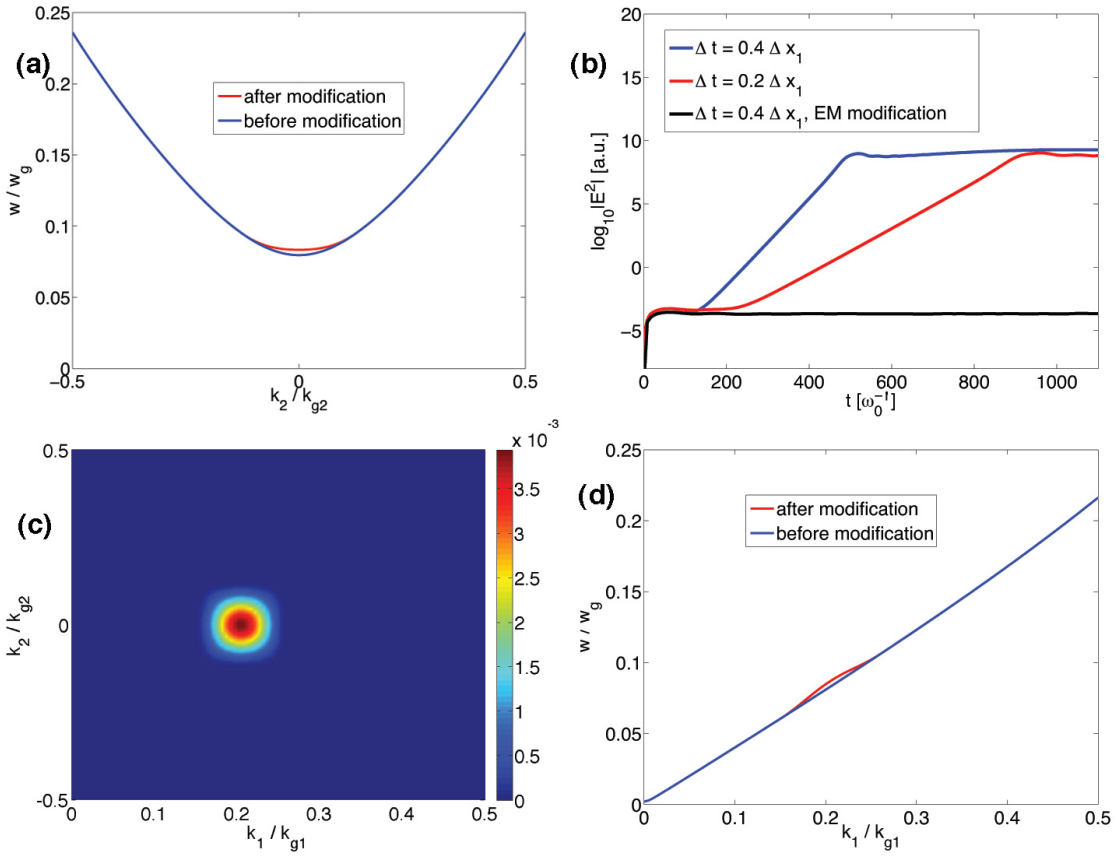

Figure 3.2: (a) shows the EM dispersion relation in vacuum before and after the modification, at line $\hat{k}_1 = 0.205$ and $\hat{k}_2 = 0$, while (c) shows the modification $|\omega - \bar{\omega}|/\omega_g$ in the fundamental Brillouin zone. (d) is the line out of (c) at $k_2 = 0$. (b) shows the E_2 energy evolution for simulations with $\Delta t = 0.4\Delta x_1$ (with and without modification), and $\Delta t = 0.2\Delta x_1$. Other simulation parameters are listed in Table 2.2.

In this section, we have shown that we can first move the $(\mu, \nu_1) = (0, 0)$ NCI modes away from physical modes by reducing the time step, then eliminate them by either applying a low-pass filter, or slightly modifying the EM dispersion in the highly localized region where the $(\mu, \nu_1) = (0, 0)$ modes reside. One can take advantage of all these strategies separately or combine them to obtain the best recipe for a particular application.

3.2 Sample simulations

In this section, we will demonstrate the approaches discussed in the previous section can efficiently mitigate the NCI in relativistic collisionless shock simulation, and in LWFA simulations in a Lorentz boosted frame, both of which involves the modeling of relativistically drifting plasma.

3.2.1 Relativistic collisionless shock

In Fig. 3.3 we present the results of two colliding plasma simulations, using the parameters in Table 3.1, with two different time steps. In these simulations we model the interaction of two counter-streaming plasma flows, each moving with a relativistic Lorentz factor of 20.0. Each plasma is initialized with a momentum distribution given by

$$f(\vec{p}) \sim \exp\left(-\frac{(p_1 - p_{10})^2}{2p_{th,1}^2}\right)\exp\left(-\frac{p_2^2}{2p_{th,2}^2}\right)\exp\left(-\frac{p_3^2}{2p_{th,3}^2}\right) \tag{3.5}$$

where p_{10} and $\vec{p}_{th}$ are listed in Table 3.1. As the two flows interpenetrate they give rise to the so-called collisionless Weibel instability [66], which slows down the flows and forms two shocks that propagate in opposite directions. In both cases we use the low-pass filter to eliminate the $(\mu, \nu_1) = (0, \pm 1)$ NCI. Comparing the $\log_{10}|B_3|$ plots in Fig. 3.3 (b) with $\Delta t = 0.4\Delta x_1$ and (c) with $\Delta t = 0.08\Delta x_1$, it is evident that when the time step is reduced, the noise originating from the NCI in the region where the two streams have not yet collided (overlap) with each other [shown in the red boxes in Fig. 3.3 (b) and (c)] is much smaller. In Fig. 3.3 (d) and (e) we also plot the FFT of the B_3 field for these same areas. The characteristic four-dot pattern of the $(\mu, \nu_1) = (0, 0)$ modes is clearly observed only for $\Delta t = 0.4\Delta x_1$. This illustrates that the $(\mu, \nu_1) = (0, 0)$ modes can limit the length of the plasma that can be simulated even if the fastest growing modes are filtered out, and that these modes can be controlled by reducing the time step. The plasma density for the smaller time step at the same physical time is shown in Fig. 3.3 (a) to show that there is no instability in the parts of the two streams that have not overlapped yet.

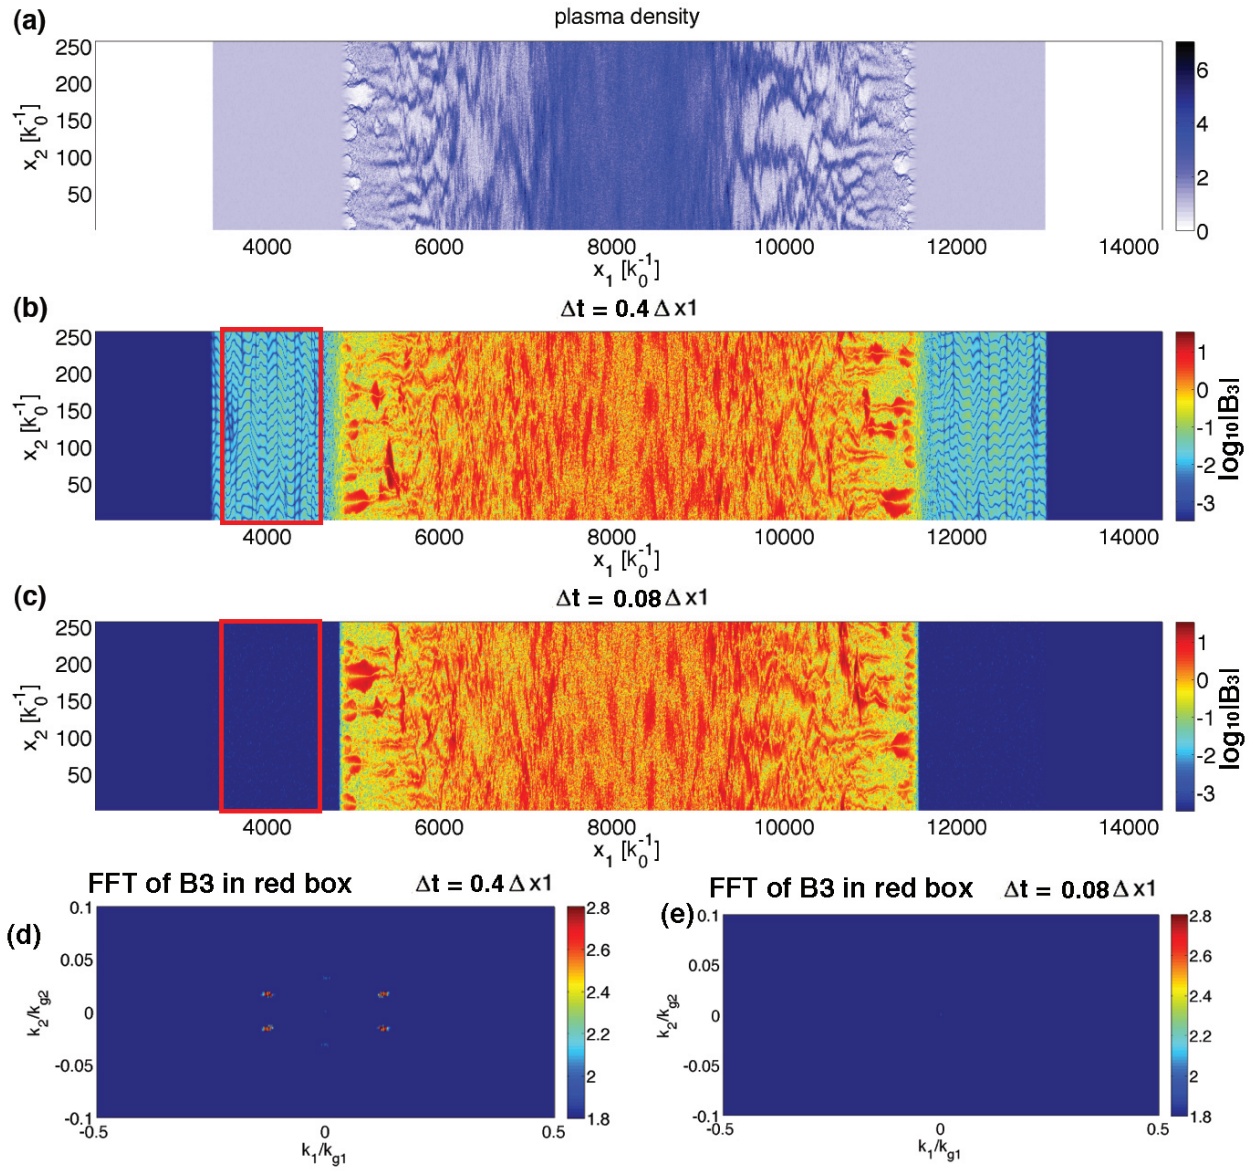

Figure 3.3: (a) shows the plasma density plot at $t = 3360\ \omega_p^{-1}$ for the $\Delta t = 0.08\Delta x_1$ case; (b) and (c) show the corresponding $\log_{10}|B_3|$ for the case $\Delta t = 0.4\Delta x_1$ and $\Delta t = 0.08\Delta x_1$, respectively. (d) and (e) shows the FFT of B_3 in the red box regions in (b) and (c), respectively.

Parameters	**Values**
grid size $(\Delta x_1, \Delta x_2)$	$(0.5k_p^{-1}, 0.5k_p^{-1})$
time step Δt	$0.4\Delta x_1$, $0.08\Delta x_1$
number of grid	32768×512
particle shape	quadratic
electron drifting momentum p_{10}	19.975 $m_e c$
electron $\vec{p}_{th}$	(0.001,0.001,0.001) $m_e c$
Ion mass ratio m_i/m_e	32

Table 3.1: Simulation parameters for the 2D shock simulation. n_p is the plasma density, and $\omega_p^2 = 4\pi q^2 n_p/m_e$, $k_p = \omega_p$ (c is normalized to 1).

3.2.2 LWFA simulation in the Lorentz boosted frame

We next present results from a LWFA boosted frame simulation in a nonlinear regime. The nonlinear regime is more challenging to simulate in the boosted frame due to self-trapping and the presence of wave harmonics. We first note that in LWFA simulations the plasma density is not really a free parameter when the simulation is done in the wakefield frame where $\gamma_b = \gamma_w \equiv \omega_0/\omega_{p0}$, ω_0 is the laser frequency. In this frame $\omega_0' = \gamma_b(\omega_0 - k_0 v_b) = 2\omega_0/\gamma_b$, and $\omega_p^2/\gamma_b = \omega_{p0}^2$ is an invariant, which leads to

$$\frac{\omega_p^2}{\gamma_b \omega_0'^2} = \frac{\omega_{p0}^2 \gamma_b^2}{4\omega_0^2} = \frac{1}{4} \tag{3.6}$$

Therefore, with respect to ω_0' the value of ω_p^2/γ_b is fixed. The time steps and cell sizes are determined with respect to ω_0', therefore ω_p^2/γ_b is not a free parameter.

In Fig. 3.4 we present results using parameters listed in Table 3.2. These parameters are the same as in Ref. [71] with $\gamma_b = 28$. It is worth noting that at the time when [71] (on which Chapter 6 of this dissertation is based) was published, we were not aware of the effect of the main NCI modes on the LWFA boosted frame simulations, and this effect is studied in more details in the simulation presented below.

The reference run used the time step $\Delta t = 0.225\Delta x_1$, and additional cases were simulated to eliminate the NCI growth: a case with a reduced time step of $\Delta t = 0.0563\Delta x_1$, and a case

with $\Delta t = 0.225\Delta x_1$ plus the EM dispersion modification (with the modification parameter $k_{1,\min}/k_{g1} = 0.151$, $k_{1,\max}/k_{g1} = 0.222$, $k_{2,\max}/k_{g2} = 0.125$, and $k_{inc,\max}/k_{g1} = 0.01$). The spatial resolution and number of simulation particles were kept fixed. In each case the low-pass filter is applied to eliminate the fastest growing NCI modes. In Figs. 3.4 (a)–(c) we show the $\log_{10}|E_2|$ for the three cases at $t = 11135\ \omega_0^{-1}$. As is shown in Fig. 3.4, the self-injected particles observed in the case of $\Delta t = 0.225\Delta x_1$ without EM dispersion modification [Fig. 3.4 (a)] are no longer seen in the case with reduced time step $\Delta t = 0.0563\Delta x_1$ [Fig. 3.4 (b)], or the $\Delta t = 0.225\Delta x_1$ case with EM dispersion modification [Fig. 3.4 (c)]. Note in the 2D OSIRIS lab frame simulation, no self-injection particles are observed.

The fact that the $\Delta t = 0.225\Delta x_1$ cases without EM dispersion modification shows self-injected particles, while in the $\Delta t = 0.225\Delta x_1$ case with dispersion modification and the $\Delta t = 0.225\Delta x_1$ case there is no self-injected particle strongly indicates the $(\mu, \nu_1) = (0, 0)$ NCI modes are interfering with the modeling of the self-injection process. Slightly modifying the EM dispersion curve does not change the accuracy of the other parts of the algorithm (e.g. Maxwell solver, pusher), therefore the only difference between the two cases is that for the modified-dispersion case there is no $(\mu, \nu_1) = (0, 0)$ NCI modes, and the absence of these unphysical modes brings the simulation results closer to the lab frame results.

As a side note, we see in the green box in Fig. 3.4 (c) there is radiation that is not seen in Fig. 3.4 (a). This is due to the fact that when we artificially create a bump in the EM dispersion relation, part of the $\vec{k}$ in the bump has a group velocity difference to the drifting velocity of the plasma larger than the $\vec{k}$ outside the bump. As a result the radiation that is in the range of these $\vec{k}$s will travel faster than the other $\vec{k}$s. We isolated the green box region and performed an FFT for the data inside the box, as shown in Figs. 3.4 (d) and (e). We can see that the range of $\vec{k}$ for the radiation that is in the front of the drifting plasma corresponds exactly to those that have a larger group velocity.

The fact that both strategies bring the boosted frame simulation results closer to the lab frame results can also be seen by transforming the on-axis wakefields E_1 back to the lab frame and comparing them with lab frame OSIRIS simulation. In Fig. 3.4 (f) we plot a lab frame time sequence of line outs of the on-axis wakefield. Here we plot the line outs

for the time steps $\Delta t = 0.225\Delta x_1$ without the EM dispersion modification (red curve), $\Delta t = 0.0563\Delta x_1$ (green curve), and $\Delta t = 0.225\Delta x_1$ with the EM dispersion modification (cyan curve), and OSIRIS lab frame data (blue curve). It shows that better agreement with the lab frame result is found for the reduced time step, and for the case with larger time step plus EM dispersion modification. For the larger time step case without the EM dispersion modification (red curve) one can see the wake is perturbed at early times before the electric field reaches its minimum value in the rear of the first bubble. This is due to the self-trapped particles which are absent for the lab frame, and the two boosted frame simulations with elimination strategies applied to eliminate the $(\mu, \nu_1) = (0, 0)$ NCI modes.

Plasma	
density n_0	$1.148 \times 10^{-3} n_0 \gamma_b$
length L	$7.07 \times 10^4 k_0^{-1}/\gamma_b$
Laser	
pulse length τ	$70.64 k_0^{-1} \gamma_b (1 + \beta_b)$
pulse waist W	$117.81 k_0^{-1}$
polarization	$\hat{3}$-direction
2D boosted frame simulation	
grid size $\Delta x_{1,2}$	$0.0982 k_0^{-1} \gamma_b (1 + \beta_b)$
time step $\Delta t/\Delta x_1$	0.225, 0.0563
number of grid ($\gamma_b = 28$)	8192×256
particle shape	quadratic

Table 3.2: Parameters for the 2D LWFA simulations, with $a_0 = 4.0$. The laser frequency ω_0 and laser wave number k_0 are used to normalize simulation parameters, and $n_0 = m_e \omega_0^2/(4\pi e^2)$.

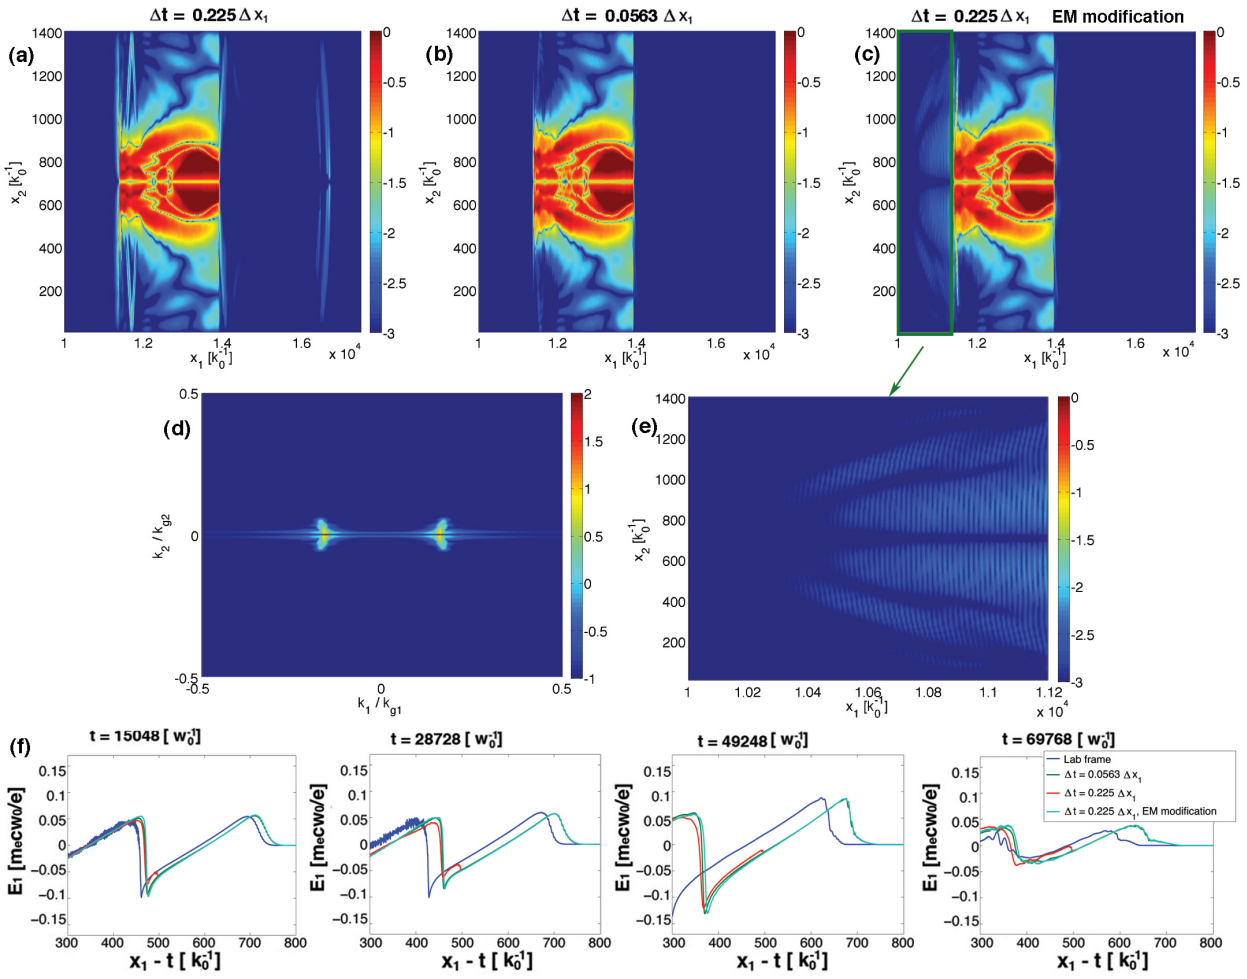

Figure 3.4: (a)–(c) shows the $\log_{10}|E_2|$ at $t = 11135\ \omega_0^{-1}$ for the cases with $\Delta t = 0.225\Delta x_1$ (with and without EM modification), and $\Delta t = 0.0563\Delta x_1$, respectively. In (e) we magnify the region in the green box in (c) to show the detailed structure of the radiation ahead of the drifting plasma, and (d) shows the corresponding FFT spectral for these radiation. (f) shows the on-axis E_1 wakefield when transforming the UPIC-EMMA simulations data back to the lab frame and compared again OSIRIS lab frame simulation data.

3.3 Summary

In this chapter, we used the theory developed in Chapter 2, and discussed in details the NCI elimination schemes for the spectral solver in which the Maxwell equations are spatially advanced in Fourier space. The main advantage of the spectral solver, comparing against the Yee solver, is that the $(\mu, \nu_1) = (0, \pm 1)$ NCI modes are highly localized in the spectral solver, and their growth rates are one order of magnitude smaller than the fastest growing $(\mu, \nu_1) = (0, \pm 1)$ modes. For the Yee solver, the $(\mu, \nu_1) = (0, \pm 1)$ and $(\mu, \nu_1) = (0, 0)$ modes have similar growth rate and the $(\mu, \nu_1) = (0, \pm 1)$ modes are not localized. In addition, when a reduced simulation time step is applied, their growth rates are reduced, and their locations move farther away from the $\vec{k} = \vec{0}$ where the physical modes reside. Furthermore, these main NCI modes in the spectral solver can be completely eliminated by creating a bump near where they are located in the $\vec{k}$ space. Meanwhile, the $(\mu, \nu_1) = (0, \pm 1)$ NCI modes located near the edge of the fundamental Brillouin zone which is similar to the Yee solver, and can be eliminated by applying a low-pass filter to the current. As for even higher order NCI modes, one can use a higher order particle shape to reduce their growth rates.

As we will see in the next chapter, the elimination scheme developed for the spectral solver provide us with a systematic approach to completely eliminate the two sets of fastest growing modes, the main, and first spatial aliasing NCI modes. While spectral solver is perfectly up to the task of NCI elimination, it does have certain shortcomings. This includes a constraint on the partition of the simulation box, which would eventually affect the scalability of the PIC code. In addition, as mentioned in Chapter 1, a spectral PIC code is relatively difficult to develop due to the complex messaging mechanism in the solver. In the next chapter, we attempt to take advantage of what we learned in this chapter, to design new Maxwell solvers that are specially designed to eliminate the NCI, while reducing the complexity in implementing them into the PIC code.

CHAPTER 4

Hybrid Yee-FFT solver for NCI elimination

In Chapter 3, we have described our approaches to eliminate the NCI based on a multi-dimensional spectral solver. This elimination scheme can efficiently eliminate the main and first spatial aliasing NCI modes, which makes it possible to perform high fidelity LWFA Lorentz boosted frame simulations, as well as colliding plasma or relativistic shock simulations. However, as mentioned in Chapter 1, it is a major effort to develop a spectral EM-PIC code can scale to millions of CPU cores, and there are only a few research groups in the world that has production spectral PIC codes for physics studies. Therefore, the following question naturally follows: is it possible to design a Maxwell solver that has similar NCI properties to that of the spectral solver, yet does not have the complexity of the spectral solver when we need to parallel it? The main motivation of this chapter is to seek answers to this question. In this chapter, we take advantage of the knowledge described in Chapter 2 and 3 as guidelines, and develop solvers that have similar NCI properties to the spectral solver, yet are much easier to implement into existing PIC codes and that may have better parallel scalability. In addition, we seek NCI elimination strategies that will work for the new quasi-3D algorithm to be discussed later.

We will first describe the design of a hybrid Yee-FFT solver, in which Maxwell's equations are solved in a hybrid (k_1, x_2, x_3) space. As we will explain in the following sections, this solver takes advantage of the fact that one only needs to apply the spectral operator in the direction the plasma is drifting, in order to obtain similar NCI properties to that of the fully spectral solver. In addition, this solver uses a similar NCI elimination strategy as that of the spectral solver. Since this solver is only spectral in one direction while retaining its FDTD structure in the other direction(s), it is much easier to implement into an existing

FDTD code compared with a multi-dimensional spectral solver. Moreover, this idea can be extended to a hybrid geometry called quasi-3D geometry, which opens the possibility of performing relativistic plasma drift simulation in quasi-3D. The quasi-3D geometry by itself offers speed ups by factors of ~ 100 for LWFA simulations.

4.1 Hybrid Yee-FFT solver

We will first describe the reasoning and the detailed algorithm of the hybrid Yee-FFT Maxwell solver. The basic idea of this solver is that the theoretical framework developed in Chapter 2 indicates that the NCI is easier to eliminate when EM waves are superluminal along the direction of the plasma drift. This can be accomplished with higher order solvers or with an FFT based solver in the drifting direction of the plasma (denoted as $\hat{1}$-direction). We note that it is more difficult to satisfy strict charge conservation (Gauss's law) for higher order finite difference solvers when using the charge conserving current deposition techniques [23]. Here we replace the finite difference operator of the first spatial derivative $\partial/\partial x_1$ in the Maxwell's equation in Yee solver with its FFT counterpart that has an accuracy greater than order N. We then correct for this change in the current deposit to maintain strict charge conservation. Without loss of generality, in the following we will briefly describe the algorithm of the Yee-FFT solver in two-dimensional (2D) Cartesian geometry. The straightforward extension to the 3D Cartesian case is also discussed.

4.1.1 Algorithm

We start from the standard algorithm for a 2D Yee solver, in which the electromagnetic fields $\vec{E}$ and $\vec{B}$ are advanced by solving Faraday's Law and Ampere's Law:

$$B^{n+\frac{1}{2}}_{1,i1,i2+\frac{1}{2}} = B^{n-\frac{1}{2}}_{1,i1,i2+\frac{1}{2}} - c\Delta t \times \frac{E^n_{3,i1,i2+1} - E^n_{3,i1,i2}}{\Delta x_2} \tag{4.1}$$

$$B^{n+\frac{1}{2}}_{2,i1+\frac{1}{2},i2} = B^{n-\frac{1}{2}}_{2,i1+\frac{1}{2},i2} + c\Delta t \times \frac{E^n_{3,i1+1,i2} - E^n_{3,i1,i2}}{\Delta x_1} \tag{4.2}$$

$$\begin{aligned} B^{n+\frac{1}{2}}_{3,i1+\frac{1}{2},i2+\frac{1}{2}} = B^{n-\frac{1}{2}}_{3,i1+\frac{1}{2},i2+\frac{1}{2}} &- c\Delta t \times \frac{E^n_{2,i1+1,i2+\frac{1}{2}} - E^n_{2,i1,i2+\frac{1}{2}}}{\Delta x_1} \\ &+ c\Delta t \times \frac{E^n_{1,i1+\frac{1}{2},i2+1} - E^n_{1,i1+\frac{1}{2},i2}}{\Delta x_2} \end{aligned} \tag{4.3}$$

$$E^{n+1}_{1,i1+\frac{1}{2},i2} = E^n_{1,i1+\frac{1}{2},i2} - 4\pi\Delta t \times j^{n+\frac{1}{2}}_{1,i1+\frac{1}{2},i2} + c\Delta t \times \frac{B^{n+\frac{1}{2}}_{3,i1+\frac{1}{2},i2+\frac{1}{2}} - B^{n+\frac{1}{2}}_{3,i1+\frac{1}{2},i2-\frac{1}{2}}}{\Delta x_2} \tag{4.4}$$

$$E^{n+1}_{2,i1,i2+\frac{1}{2}} = E^n_{2,i1,i2+\frac{1}{2}} - 4\pi\Delta t \times j^{n+\frac{1}{2}}_{2,i1,i2+\frac{1}{2}} - c\Delta t \times \frac{B^{n+\frac{1}{2}}_{3,i1+\frac{1}{2},i2+\frac{1}{2}} - B^{n+\frac{1}{2}}_{3,i1-\frac{1}{2},i2+\frac{1}{2}}}{\Delta x_1} \tag{4.5}$$

$$\begin{aligned} E^{n+1}_{3,i1,i2} = E^n_{3,i1,i2} - 4\pi\Delta t \times j^{n+\frac{1}{2}}_{3,i1,i2} &+ c\Delta t \times \frac{B^{n+\frac{1}{2}}_{2,i1+\frac{1}{2},i2} - B^{n+\frac{1}{2}}_{2,i1-\frac{1}{2},i2}}{\Delta x_1} \\ &- c\Delta t \times \frac{B^{n+\frac{1}{2}}_{1,i1,i2+\frac{1}{2}} - B^{n+\frac{1}{2}}_{1,i1,i2-\frac{1}{2}}}{\Delta x_2} \end{aligned} \tag{4.6}$$

where the EM fields $\vec{E}$ and $\vec{B}$, and current $\vec{j}$ are defined with the proper half-grid offsets according to the Yee mesh [21]. If we perform a Fourier transform of Eq. (4.1)–(4.6) in both x_1 and x_2, and in time, Maxwell's equations reduce to

$$[\omega]\vec{B} = c[\vec{k}] \times \vec{E} \tag{4.7}$$

$$[\omega]\vec{E} = -c[\vec{k}] \times \vec{B} - 4\pi i\vec{j} \tag{4.8}$$

where

$$[\vec{k}] = \left(\frac{\sin(k_1\Delta x_1/2)}{\Delta x_1/2}, \frac{\sin(k_2\Delta x_2/2)}{\Delta x_2/2}, 0\right) \qquad [\omega] = \frac{\sin(\omega\Delta t/2)}{\Delta t/2} \tag{4.9}$$

In vacuum where $\vec{j} = 0$, the corresponding numerical dispersion relation for the EM waves is

$$[\omega]^2 = c^2([k]_1^2 + [k]_2^2) \tag{4.10}$$

The idea of a hybrid Yee-FFT solver is to keep the finite difference operator

$$[k]_2 = \sin(k_2\Delta x_2/2)/(\Delta x_2/2) \tag{4.11}$$

in the directions transverse to the drifting direction, while replacing the finite difference operator $[k]_1$ in the drifting direction with its spectral counterpart $[k]_1 = k_1$. To achieve this, in the hybrid solver we will solve Maxwell's equations in k_1 space. The current is deposited locally using a rigorous charge conserving scheme [23] that is valid for second order finite difference operator. For the EM field and current, we first perform an FFT along x_1 so that all fields are defined in (k_1, x_2) space. After that we apply a correction to the current in the drifting direction

$$\tilde{j}_1^{n+\frac{1}{2}} = \frac{\sin(k_1\Delta x_1/2)}{k_1\Delta x_1/2} j_1^{n+\frac{1}{2}} \tag{4.12}$$

where $\tilde{j}_1$ is the corrected current. In [54], the current is also corrected where they combine a pure FFT solver with a charge conserving current deposit. This correction ensures that Gauss's Law is satisfied throughout the duration of the simulation if it is satisfied initially, as will be discussed in more detail in section 4.1.3. After the current correction, we advance the EM field as

$$B^{n+\frac{1}{2}}_{1,\kappa 1,i2+\frac{1}{2}} = B^{n-\frac{1}{2}}_{1,\kappa 1,i2+\frac{1}{2}} - c\Delta t \times \frac{E^n_{3,\kappa 1,i2+1} - E^n_{3,\kappa 1,i2}}{\Delta x_2} \tag{4.13}$$

$$B^{n+\frac{1}{2}}_{2,\kappa 1,i2} = B^{n-\frac{1}{2}}_{2,\kappa 1,i2} - i\xi^+ k_1 c\Delta t E^n_{3,\kappa 1,i2} \tag{4.14}$$

$$B^{n+\frac{1}{2}}_{3,\kappa 1,i2+\frac{1}{2}} = B^{n-\frac{1}{2}}_{3,\kappa 1,i2+\frac{1}{2}} + i\xi^+ k_1 c\Delta t E^n_{2,\kappa 1,i2+\frac{1}{2}} + c\Delta t \times \frac{E^n_{1,\kappa 1,i2+1} - E^n_{1,\kappa 1,i2}}{\Delta x_2} \tag{4.15}$$

$$E^{n+1}_{1,\kappa 1,i2} = E^n_{1,\kappa 1,i2} - 4\pi\Delta t \times \tilde{j}^{n+\frac{1}{2}}_{1,\kappa 1,i2} + c\Delta t \times \frac{B^{n+\frac{1}{2}}_{3,\kappa 1,i2+\frac{1}{2}} - B^n_{3,\kappa 1,i2-\frac{1}{2}}}{\Delta x_2} \tag{4.16}$$

$$E^{n+1}_{2,\kappa 1,i2+\frac{1}{2}} = E^n_{2,\kappa 1,i2+\frac{1}{2}} - 4\pi\Delta t \times j^{n+\frac{1}{2}}_{2,\kappa 1,i2+\frac{1}{2}} + i\xi^- k_1 c\Delta t B^{n+\frac{1}{2}}_{3,\kappa 1,i2+\frac{1}{2}} \tag{4.17}$$

$$\begin{aligned} E^{n+1}_{3,\kappa 1,i2} = {} & E^n_{3,\kappa 1,i2} - 4\pi\Delta t \times j^{n+\frac{1}{2}}_{3,\kappa 1,i2} - i\xi^- k_1 c\Delta t B^{n+\frac{1}{2}}_{2,\kappa 1,i2} \\ & - c\Delta t \times \frac{B^n_{1,\kappa 1,i2+\frac{1}{2}} - B^n_{1,\kappa 1,i2-\frac{1}{2}}}{\Delta x_2} \end{aligned} \tag{4.18}$$

where $k_1 = 2\pi\kappa_1/N$ and N is the number of grids in x_1 direction, and $\kappa_1 = 0, 1, \ldots, N/2-1$ is the mode number. Note in the hybrid solver, the EM fields $\vec{E}$, $\vec{B}$, and current $\vec{j}$ have the

same temporal and spatial centering as in the Yee solver, and

$$\xi^{\pm} = \exp\left(\pm\frac{k_1\Delta x_1}{2}i\right) \tag{4.19}$$

is the phase shifting due to the half grid offsets of the E_1, $B_{2,3}$, and j_1 in the $\hat{1}$-direction. Compared with the standard Yee solver algorithm, it is evident that if we replace $-ik_1$ with the corresponding finite difference form we can recover the standard 2D Yee algorithm.

4.1.2 Courant condition

The Courant condition of the hybrid solver can be easily derived from the corresponding numerical EM dispersion Eq. (4.10). Substituting into Eq. (4.10) the finite difference operator in time $[\omega]$

$$[\omega] = \frac{\sin(\omega\Delta t/2)}{\Delta t/2} \tag{4.20}$$

and the finite difference operators in space

$$[k]_1 = k_1 \qquad [k]_2 = \frac{\sin(k_2\Delta x_2/2)}{\Delta x_2/2} \tag{4.21}$$

we can obtain the corresponding constraint on the time step

$$\frac{\Delta t}{2}\sqrt{k_1^2 + \left(\frac{\sin(k_2\Delta x_2/2)}{\Delta x_2/2}\right)^2} \leq 1 \tag{4.22}$$

Note the $\vec{k}$ range of the fundamental Brillouin zone is $|k_1| \leq \pi/\Delta x_1$, $|k_2| \leq \pi/\Delta x_2$, so we can obtain the Courant limit for the hybrid solver

$$\Delta t \leq \frac{2}{\sqrt{\frac{\pi^2}{\Delta x_1^2} + \frac{4}{\Delta x_2^2}}} \tag{4.23}$$

For square cells with $\Delta x_1 = \Delta x_2$, this reduces to $\Delta t \leq 0.537\Delta x_1$.

4.1.3 Charge conservation

In the hybrid Yee-FFT solver, we rely on the Faraday's Law and Ampere's Law to advance the EM field. On the other hand, the local charge conserving current deposition [23] ensures

the second-order-accurate finite difference representation of the continuity equation,

$$\overline{\frac{\partial}{\partial t}}\rho^n_{i1,i2} + \frac{j^{n+\frac{1}{2}}_{1,i1+\frac{1}{2},i2} - j^{n+\frac{1}{2}}_{1,i1-\frac{1}{2},i2}}{\Delta x_1} + \frac{j^{n+\frac{1}{2}}_{2,i1,i2+\frac{1}{2}} - j^{n+\frac{1}{2}}_{2,i1,i2-\frac{1}{2}}}{\Delta x_2} = 0 \tag{4.24}$$

is satisfied, where

$$\overline{\frac{\partial}{\partial t}}G^n = \frac{G^{n+1} - G^n}{\Delta t} \tag{4.25}$$

where G^n is an arbitrary scalar quantity. Therefore, when combining this scheme with the second order accurate Yee solver, Gauss's Law is rigorously satisfied at every time step if it is satisfied at $t = 0$. However, when the hybrid solver is used together with the charge conserving current deposition scheme, we need to apply a correction to the current, as shown in Eq. (4.12), in order that the Gauss's Law is satisfied at every time step. This can be seen by first performing a Fourier transform in the x_1 direction for Eq. (7.9),

$$\overline{\frac{\partial}{\partial t}}\rho^n_{\kappa 1,i2} - i\frac{\sin(k_1\Delta x_1/2)}{\Delta x_1/2} j^{n+\frac{1}{2}}_{\kappa 1,i2} + \frac{j^{n+\frac{1}{2}}_{2,\kappa 1,i2+\frac{1}{2}} - j^{n+\frac{1}{2}}_{2,\kappa 1,i2-\frac{1}{2}}}{\Delta x_2} = 0 \tag{4.26}$$

then applying the divergence operator of the hybrid solver to the left and right hand side of the Ampere's Law, Eq. (4.16)–(4.18). Using Eq. (4.26), we obtain

$$\overline{\frac{\partial}{\partial t}}\left(-4\pi\rho^n_{\kappa 1,i2} - ik_1 E^n_{1,\kappa 1,i2} + \frac{E^n_{2,\kappa 1,i2+\frac{1}{2}} - E^n_{2,\kappa 1,i2-\frac{1}{2}}}{\Delta x_2}\right) = 0 \tag{4.27}$$

which shows that if Gauss's Law, which for the 2D hybrid solver is given by,

$$-ik_1 E^n_{1,\kappa 1,i2} + \frac{E^n_{2,\kappa 1,i2+\frac{1}{2}} - E^n_{2,\kappa 1,i2-\frac{1}{2}}}{\Delta x_2} = 4\pi\rho^n_{\kappa 1,i2} \tag{4.28}$$

is satisfied at $t = 0$, it is satisfied at each time step. We note that this correction can effectively broaden the range of cells over which a particle contributes to current. This is discussed in more details in the next chapter.

4.1.4 3D Cartesian geometry

It is straightforward to extend the hybrid solver to 3D cartesian geometry. In 3D Cartesian coordinates, we solve Maxwell's equation in (k_1, x_2, x_3) space where we use the same second

order accurate finite difference form of the Yee solver in the $\hat{2}$ and $\hat{3}$ directions. As in the 2D Cartesian case, the current correction is applied to j_1 to ensure that Gauss's Law is satisfied at each time step. We have implemented the hybrid solver in 2D and 3D with current correction in our finite-difference-time-domain (FDTD) code OSIRIS [30].

4.2 NCI in hybrid solver

To investigate the NCI properties of the hybrid solver, we first consider its corresponding numerical dispersion relation. Employing the general theoretical framework established in Chapter 2, we can calculate in detail the NCI modes for any Maxwell solver. The roots of the numerical dispersion relation that lead to the NCI can be found numerically by solving Eq. (2.18), or by the analytical expression in Eqs. (2.36) and (2.37). For the Yee solver the k space representation of the finite difference operator is

$$[k]_i = \frac{\sin(k_i \Delta x_i/2)}{\Delta x_i/2} \tag{4.29}$$

where $i = 1, 2$ in 2D. Meanwhile, in the hybrid solver the $\vec{k}$ space operator in the drifting direction is replaced with that of the spectral solver $[k]_1 \to k_1$. By substituting the respective operators for each direction into Eqs. (2.36) and (2.37), we can rapidly find the set of NCI modes for the hybrid solver. In Fig. 4.1 (a)–(d), we plot the $(\mu, \nu_1) = (0, 0)$ and $(\mu, \nu_1) = (0, \pm 1)$ modes for the hybrid and spectral solvers by scanning over the (k_1, k_2) space in the fundamental Brillouin zone and solve for the growth rates of the corresponding unstable modes. The parameters used to generate this plot are listed in Table 4.1.

We can see from Fig. 4.1 (a) and (b) that the $(\mu, \nu_1) = (0, \pm 1)$ NCI modes of the two solvers reside near the edge of the fundamental Brillouin zone, although the patterns are slightly different due to their different finite difference operators in the $\hat{2}$-direction, which leads to the slightly different EM dispersion curves. In Fig. 4.1 (e) and (f) we show how different EM dispersion curves leads to different $(\mu, \nu_1) = (0, \pm 1)$ NCI modes for the two solvers. These modes are distinct, and far removed from the modes of physical interest, and are relatively easy to eliminate.

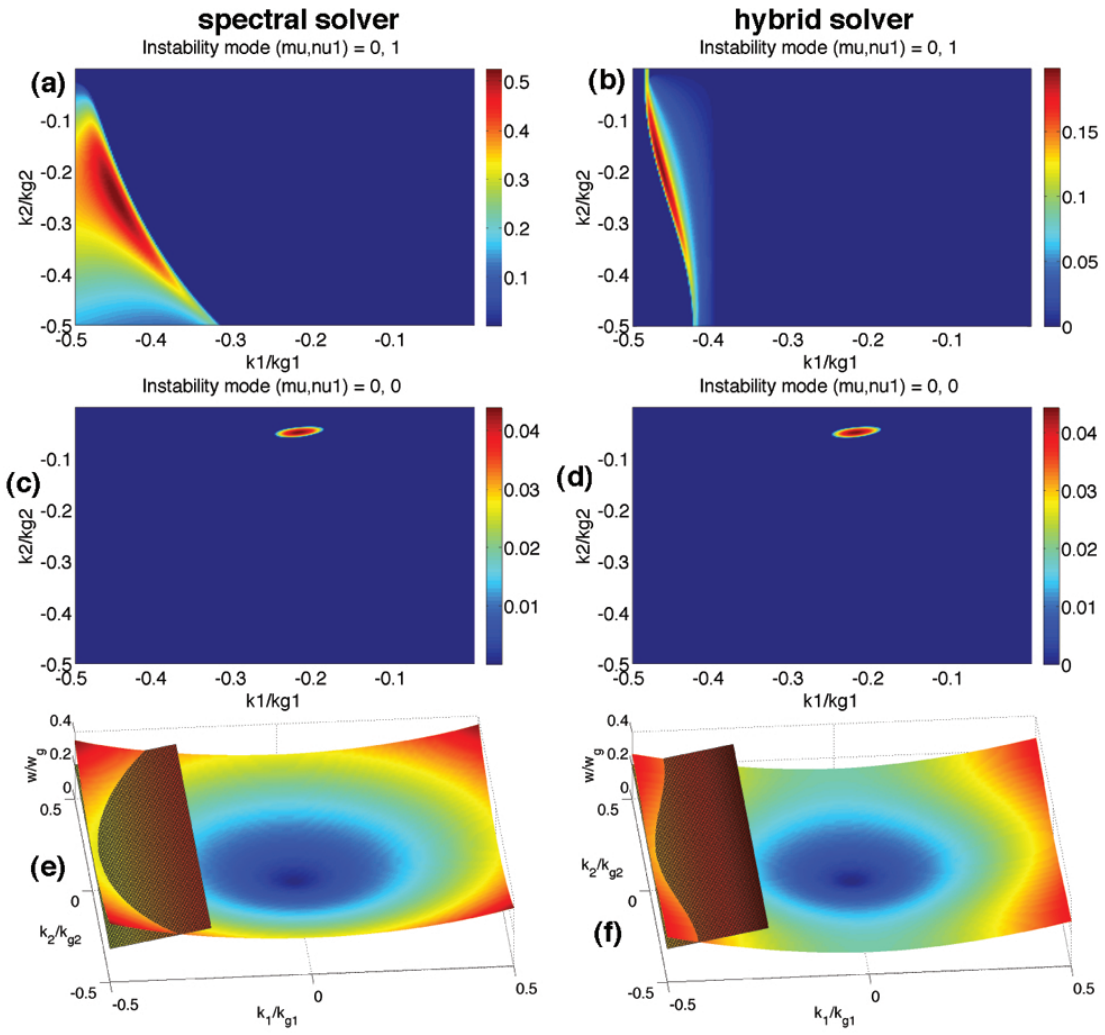

Figure 4.1: The pattern of the $(\mu, \nu_1) = (0, \pm 1)$ modes for the two solvers are shown in (a) and (b). The pattern of the $(\mu, \nu_1) = (0, 0)$ modes for two solvers are shown in (c) and (d). The intersection between the EM dispersion relations with the first spatial aliasing beam modes for the full spectral solver and the hybrid solver are shown in (e) and (f). When generating these plots we use $\Delta x_1 = \Delta x_2 = 0.2\ k_p^{-1}$, and $\Delta t = 0.08\ \omega_p^{-1}$. Other parameters are listed in Table 4.1.

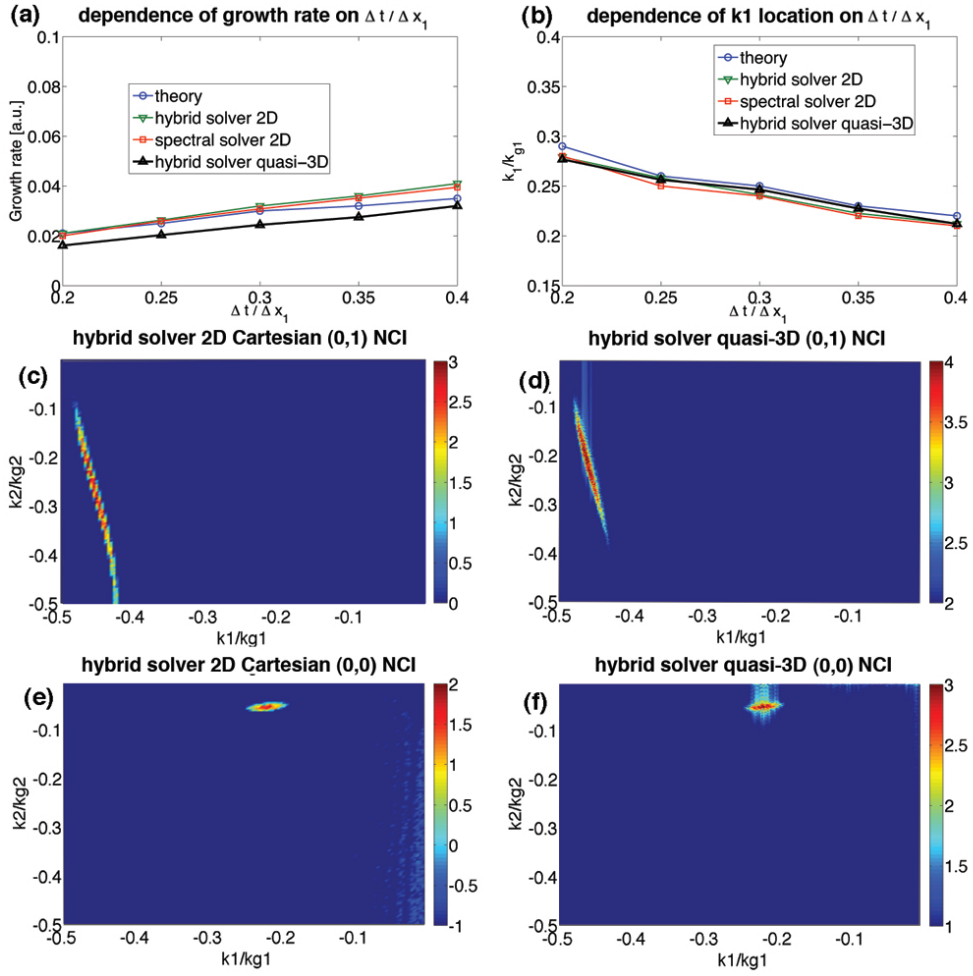

Figure 4.2: In (a) and (b) the dependence of the growth rate and k_1 for the fastest growing $(\mu, \nu_1) = (0, 0)$ mode on the time step is shown. The four lines correspond to the theoretical prediction for the hybrid solver in 2D, results from OSIRIS and UPIC-EMMA simulations for the spectral and hybrid solvers in 2D Cartesian geometry, and results for the hybrid solver in the quasi-3D geometry (where the k_2 is obtained from a Hankel transform). In (c)–(f) the spectrum of E_2 (E_ρ) is plotted for OSIRIS simulations with the hybrid solver in 2D Cartesian or the quasi-3D geometry. In (c) and (d) results from runs where no filter in k_1 is used to eliminate the $(\mu, \nu_1) = (0, \pm 1)$ modes. In (e) and (f) a filter in k_1 was used to eliminate the $(\mu, \nu_1) = (0, \pm 1)$ modes and now the $(\mu, \nu_1) = (0, 0)$ modes are seen.

Parameters	**Values**
grid size $(k_p\Delta x_1, k_p\Delta x_2)$	$(0.2, 0.2)$
time step $\omega_p\Delta t$	$0.4k_p\Delta x_1$
boundary condition	Periodic
simulation box size (k_pL_1, k_pL_2)	51.2×51.2
plasma drifting Lorentz factor	$\gamma = 50.0$
plasma density	$n/n_p = 100.0$

Table 4.1: Simulation parameters for the 2D relativistic plasma drift simulation. n_p is the reference plasma density, and $\omega_p^2 = 4\pi q^2 n_p/m_e$, $k_p = \omega_p$ (c is normalized to 1).

More importantly, we see from Fig. 4.1 (c) and (d) that the hybrid solver leads to $(\mu, \nu_1) = (0, 0)$ NCI modes that are very similar to their spectral solver counterpart. The pattern of the $(\mu, \nu_1) = (0, 0)$ modes for these two solvers are both four dots (in 2D) and highly localized in the fundamental Brillouin zone. We also use the theory to perform parameter scans to study the dependence of growth rates (of the fastest growing mode) and the locations in k space of the $(\mu, \nu_1) = (0, 0)$ modes on $\Delta t/\Delta x_1$ for the hybrid solver, and compare this result against that of the fully spectral solver, as shown in Fig. 4.2 (a) and (b). We likewise carried out OSIRIS simulations using the hybrid solver and UPIC-EMMA [71, 58] using the spectral solver, to compare against theoretical results. Very good agreement is found between theory and simulations. Fig. 4.2 (a) and (b) show that both the k_1 location, and growth rates of the $(\mu, \nu_1) = (0, 0)$ modes are almost identical for the two solvers. This indicates that, just like the spectral solver, the growth rate of the $(\mu, \nu_1) = (0, 0)$ modes of the hybrid solver is reduced, while their location in k_1 increases, when the time step is reduced.

In Fig. 4.2 (c) and (e) we show the locations of the unstable $(\mu, \nu_1) = (0, \pm 1)$, and $(\mu, \nu_1) = (0, 0)$ NCI modes for the hybrid solver in OSIRIS for 2D Cartesian geometry. The agreement between Fig. 4.2 (c) and Fig. 4.1 (b), and between Fig. 4.2 (e) and Fig. 4.1 (d) are excellent.

The main advantage of the purely spectral solver regarding its NCI properties in com-

parison to a purely FDTD solver is that the superluminal dispersion relation makes it much easier to eliminate the NCI modes at $(\mu, \nu_1) = (0, 0)$: the modes have a growth rate that is one order of magnitude smaller than that for the $(\mu, \nu_1) = (0, \pm 1)$ modes, their locations are highly localized in $\vec{k}$ space, and they can be moved away from the modes of physical interest by reducing the time step. We showed above that similar NCI properties can be achieved by using a hybrid FDTD-spectral solver, where the Maxwell's equation are solved in Fourier space only in the direction of the plasma drift. Comparing with an EM-PIC code using a multi-dimensional spectral solver which solves Maxwell's equation in $\vec{k}$ space, there are advantages when solving it in (k_1, x_2) space in 2D [and (k_1, x_2, x_3) space in 3D]. Firstly, the hybrid solver saves the FFT in the other directions; secondly, since the solver is FDTD in the directions transverse to the drifting direction, it is easier to integrate the algorithm into existing FDTD codes such as OSIRIS where the parallelizations and boundary conditions in the transverse direction can remain untouched. Last but perhaps most important, the idea that one can obtain preferable NCI properties by solving Maxwell's equation in k_1 space in the drifting direction can be readily extended to the quasi-3D algorithm [51], as we can solve the Maxwell's equation in (k_1, r, ϕ) space.

4.3 Elimination of the NCI modes

In Chapter 3, we proposed strategies to eliminate the NCI in the spectral solver. These strategies can be readily applied to the hybrid solver. For square (or cubic) cells, the pattern of the fastest growing modes resides in a narrow range of k_1 near the edge of the fundamental Brillouin zone. Therefore we can apply a low-pass filter in k_1 to the current to eliminate the fastest growing modes. Since the fields are already in k_1 space when solving the Maxwell's equations, the filtering can be done efficiently by applying a form factor to the current only in k_1.

As for the $(\mu, \nu_1) = (0, 0)$ mode, if they are near the main or higher order harmonics of the physical modes, we can move them away towards higher $|k_1|$, and reduce their growth rates by simply reducing the time step. To further mitigate the $(\mu, \nu_1) = (0, 0)$ NCI modes,

when they are far away from the physical modes, one can modify the EM dispersion relation to completely eliminate them. In Fig. 4.3 we plot how the modification is accomplished in the hybrid solver. As shown in Fig. 4.3 (a) except for the bump region, for most k_1 the $[k]_1$ for a particular k_1 is k_1 itself; near the bump, the $[k]_1$ for k_1 is $k_1 + \Delta k_{mod}$, where Δk_{mod} is a function of k_1 with

$$\Delta k_{mod} = \Delta k_{mod,\max} \cos\left(\frac{k_1 - k_{1m}}{k_{1l} - k_{1m}} \frac{\pi}{2}\right)^2 \tag{4.30}$$

where k_{1l}, k_{1u} are the lower and upper k_1 to be modified, $k_{1m} = (k_{1l} + k_{1u})/2$, and $\Delta k_{mod,\max}$ is the maximum value of Δk_{mod}. The values of k_{1l}, k_{1u} and $\Delta k_{mod,\max}$ are determined by the position of the $(\mu, \nu_1) = (0, 0)$ modes and their growth rates. According to the NCI theory, for the parameters in Table 4.1, when the $[k]_1$ is as defined in Fig. 4.3 (a) (with $k_{1l}/k_{g1} = 0.15$, $k_{1u}/k_{g1} = 0.26$, and $\Delta k_{mod,\max}/k_{g1} = 0.01$), there is no unstable $(\mu, \nu_1) = (0, 0)$ NCI modes, i.e., the $(\mu, \nu_1) = (0, 0)$ mode has a theoretical growth rate of zero. To verify the theoretical results in the hybrid solver, in Fig. 4.3 (b) we plot the E_2 energy growth with and without the modification. In these simulations we used the parameters in Table 4.1. The blue curve in Fig. 4.3 (b) represents the case without the modification, while the red and black curves are those with the modification to k_1. The cases with blue and red curves used quadratic particle shapes, while the case for the black curve used cubic particle shapes. We have likewise plotted the E_2 spectra at the time point $t = 3200\ \omega_0^{-1}$ indicated in Fig. 4.3 (c) and (d) for the two cases with the modifications (red and black curves in 4.3 (b)). We can see from Fig. 4.3 (b) and (c) that after the modification, the growth rate of the $(\mu, \nu_1) = (0, 0)$ NCI modes reduces to zero. Meanwhile, the red curve rises later in time due to the $(\mu, \nu_1) = (\pm 1, \pm 2)$ NCI modes. As we showed in Chapter 2 the growth rate of these higher order modes can be reduced by using higher order particle shape. Therefore when cubic particle shapes are used, as is the case for the black curve, the $(\mu, \nu_1) = (\pm 1, \pm 2)$ NCI modes do not grow exponentially and are therefore much less observable in the corresponding spectrum at $t = 3200\ \omega_0^{-1}$ in Fig. 4.3 (d) as compared to 4.3 (c).

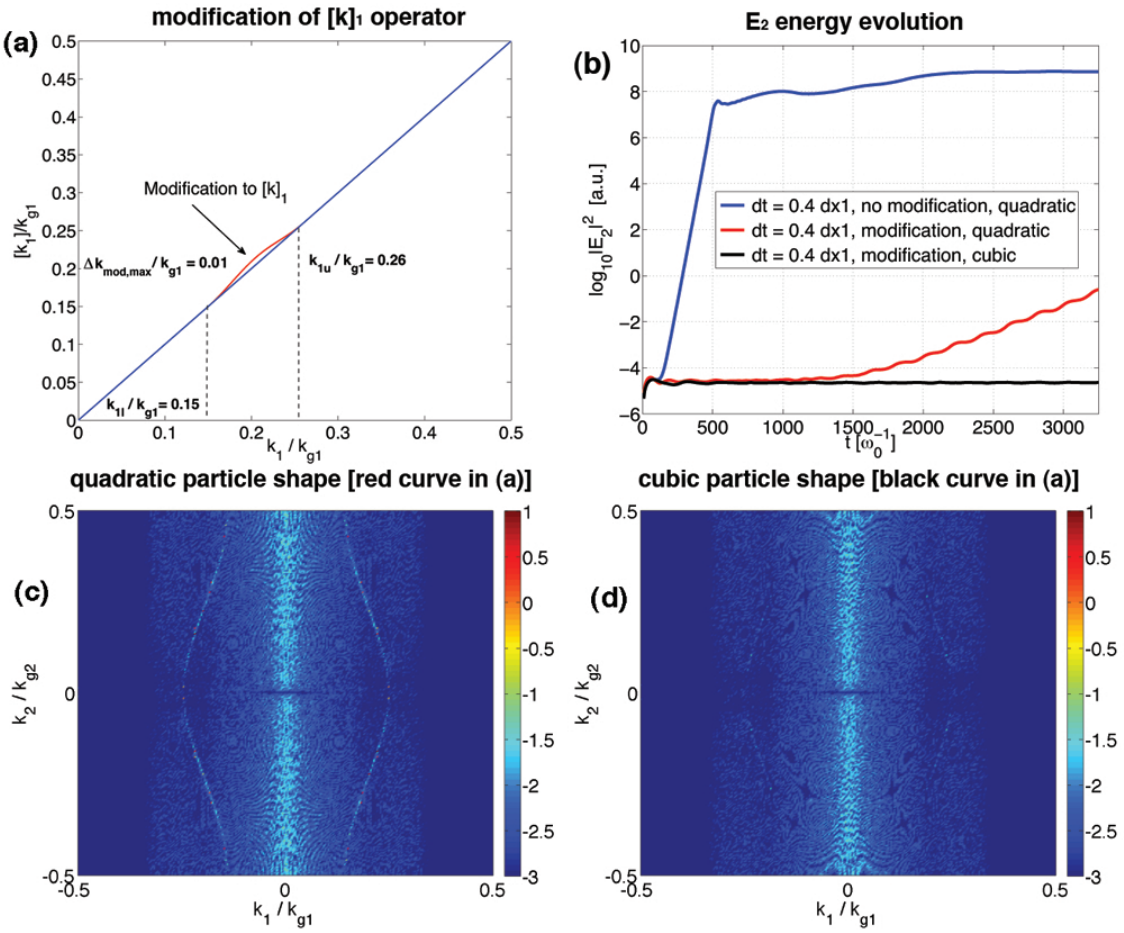

Figure 4.3: In (a) the perturbation to $[k]_1$ that is used to eliminate the $(\mu, \nu_1) = (0, 0)$ NCI modes is shown. In (b) the evolution of the $\log_{10} |E_2|^2$ for a reference case and for two cases with the EM dispersion modification (one with quadratic and another with cubic particle shapes). In (c) and (d), the spectrum of E_2 at $t = 3200\ \omega_0^{-1}$ is shown for the two cases with the EM dispersion modifications. In (c) quadratic particle shapes are used, while in (d) cubic particle shapes are used.

4.4 hybrid solver in quasi-3D algorithm

As mentioned at the beginning of this chapter, the idea of the hybrid solver can be easily incorporated into the quasi-3D algorithm [51, 55] in which the fields and current are expanded into azimuthal Fourier modes. We can obtain the hybrid Yee-FFT solver for the quasi-3D algorithm by using FFTs in the $\hat{z}$ (or $\hat{x}_1$) direction and finite difference operators in $\hat{r}$ (or $\hat{x}_2$) direction in the equations for each azimuthal mode. Note in quasi-3D OSIRIS we use a charge conserving current deposition scheme for the Yee solver (as described in [55]), therefore for the hybrid solver adapted for the quasi-3D algorithm we can apply the same current correction for the use of FFTs to j_z in order that the Gauss's Law is satisfied throughout the duration of the simulation.

The NCI properties of the hybrid solver for the quasi-3D algorithm are similar to that of the 2D Cartesian geometry [57]. While a rigorous NCI theory for the quasi-3D algorithm is still under development, we can empirically investigate the NCI for this geometry through simulation. In Fig. 4.2 (d) and (f) we plot the E_r data at a time during the exponential growth of the EM fields due to the NCI, which shows the $(\mu, \nu_1) = (0, \pm 1)$ and $(\mu, \nu_1) = (0, 0)$ modes for the hybrid solver in quasi-3D geometry. For the E_r data, we conduct an FFT in x_1 and a Hankel transform in x_2. Similarly to the 2D Cartesian case, we isolate the $(\mu, \nu_1) = (0, 0)$ modes by applying a low-pass filter in the current in k_1 space to eliminate the fastest growing $(\mu, \nu_1) = (0, \pm 1)$ NCI modes. The parameters used in the simulations are listed in Table 4.1, and a conducting boundary is used for the upper r boundary. We kept azimuthal modes of $m = -1, 0, 1$ in the simulations.

By comparing Figs. 4.2 (c)–(f) it can be seen that the pattern of the NCI modes are similar for the (x_2, x_1) and (r, z) geometries. We have also plotted the dependence of the growth rate and k_1 position of the $(\mu, \nu_1) = (0, 0)$ NCI modes for the quasi-3D geometry in Fig. 4.2 (a) and (b). These plots show that when the time step decreases the growth rates of the $(\mu, \nu_1) = (0, 0)$ NCI modes in the quasi-3D geometry decreases, while the k_1 position increases (and move away from the physical modes), in a nearly similar fashion to 2D Cartesian geometry. This indicates that the same strategies for eliminating NCI in 2D

Cartesian geometry can be applied to the quasi-3D geometry. The fastest growing modes residing at the edge of the fundamental Brillouin zone can be eliminated by applying a low-pass filter in the current. The $(\mu, \nu_1) = (0, 0)$ NCI modes can be mitigated by either reducing the time step to lower the growth rate and move the modes away from the physics in k_1 space, or by modifying the $[k]_1$ operator as discussed in section 4.3 to create a bump in the EM dispersion along the k_1 direction. We have implemented the modification to the $[k]_1$ operator into the hybrid solver for the quasi-3D OSIRIS code, and have confirmed that this modification completely eliminate the $(\mu, \nu_1) = (0, 0)$ NCI modes. The coefficients used for the modification are the same as those for the 2D Cartesian case discussed in section 4.3.

4.5 Sample simulations

In this section, we present preliminary results of Lorentz boosted frame LWFA simulations using the hybrid solver in OSIRIS. In these simulations, to eliminate the first spatial aliasing NCI modes, we use a low-pass filtering for the current of the form

$$F(k_1) = \begin{cases} 1, & |k_1| < f_l k_{g1} \\ \sin^2\left(\frac{k_1 - f_u k_{g1}}{f_l k_{g1} - f_u k_{g1}} \frac{\pi}{2}\right), & f_l k_{g1} \le k_1 \le f_u k_{g1} \\ 0, & |k_1| > f_u k_{g1} \end{cases} \tag{4.31}$$

The filter retains the k_1 modes smaller than $f_l k_{g1}$ and cuts off the modes larger than $f_u k_{g1}$. A $\sin^2$ function is used between $f_l k_{g1}$ and $f_u k_{g1}$ for smooth connection between unity and zero.

For comparison, we performed simulations with the same parameters using UPIC-EMMA which uses a spectral Maxwell solver. Table 4.2 lists the simulation parameters. We use a moving antenna in both cases to launch lasers into the plasma. The results are summarized in Fig. 4.4.

In Fig. 4.4 (a)–(b) the E_1 field at $t' = 3955\omega_0^{-1}$ for simulations with both the hybrid solver and spectral solver in the Lorentz boosted frame are plotted, where ω_0 is the laser frequency in the lab frame. Both the spectral solver and hybrid solver give similar boosted

frame results, and there is no evidence of NCI affecting the physics in either case. We plot the line out of the on-axis wakefield in Fig. 4.4 (c), which shows very good agreement with one another. The very good agreement can also be seen when we transformed the boosted frame data back to the lab frame. In Fig. 4.4 (d)–(h) we plot the on-axis E_1 field for the OSIRIS lab frame data, the transformed data for the OSIRIS boosted frame simulation with the hybrid solver, and the transformed data from the UPIC-EMMA boosted simulation at several values of time in the lab frame. As seen in Fig. 4.4 (d)–(h), the transformed data from the two boosted frame simulations agrees very well with each other.

In this chapter, we mainly focus on LWFA simulations in a Lorentz boosted frame. However, it is worth pointing out that the hybrid solver can likewise be used for LWFA lab frame simulations with a moving window. When a self-injected or externally injected electron beam is accelerated by the wakefield, it will also suffer from numerical Cerenkov radiation (NCR), and may even be susceptible to the NCI in some cases. The resulting unphysical EM fields can lead to unphysical emittance growth. Applying the hybrid solver in lab frame simulations will greatly reduce the NCR, which may lead to more accurate emittance values. As a result, although not shown in this chapter, we likewise benchmarked the hybrid solver with Yee solver in LWFA lab frame simulation by comparing the wake fields and laser evolution in the two cases, and very good agreement was obtained. An example of how the hybrid solver in the lab frame can limit unphysical emittance growth is shown in Appendix B.

4.6 Summary

In this chapter we proposed to use a hybrid Yee-FFT and a rigorous charge conserving current deposit for solving Maxwell's equations in order to eliminate the numerical Cerenkov instability in PIC codes when modeling plasmas or beams that drift with relativistic speeds in a particular direction. In this solver we solve the Maxwell's equation in k_1 space along the drifting direction ($\hat{x}_1$ direction), and use second order finite difference representation for the derivatives in the other directions. This provides greater than N-th order accuracy for

Plasma	
density n_0	$1.148 \times 10^{-3} n_0 \gamma_b$
length L	$7.07 \times 10^4 k_0^{-1}/\gamma_b$
Laser	
pulse length τ	$70.64 k_0^{-1} \gamma_b (1+\beta_b)$
pulse waist W	$117.81 k_0^{-1}$
polarization	$\hat{3}$-direction
normalized vector potential a_0	4.0
Lab frame simulation ($\gamma_b = 1$)	
grid size $(\Delta x_1, \Delta x_2)$	$(0.2 k_0^{-1}, 2.75 k_0^{-1})$
time step $\Delta t/\Delta x_1$	0.996
number of grid (moving window)	4000×512
particle shape	quadratic
2D boosted frame simulation ($\gamma_b = 14$)	
grid size $\Delta x_{1,2}$	$0.0982 k_0^{-1} \gamma_b (1+\beta_b)$
time step $\Delta t/\Delta x_1$	0.125
number of grid	8192×512
particle shape	quadratic
low-pass filter (f_l, f_u)	(0.265,0.31)

Table 4.2: Parameters for a 2D LWFA simulations in the lab frame and Lorentz boosted frame that were used for in 2D Cartesian geometry with the hybrid solver in OSIRIS and with a fully spectral solver in UPIC-EMMA. The laser frequency ω_0 and number k_0 in the lab frame are used to normalize simulation parameters. The density is normalized to the critical density in the lab frame, $n_0 = m_e \omega_0^2/(4\pi e^2)$.

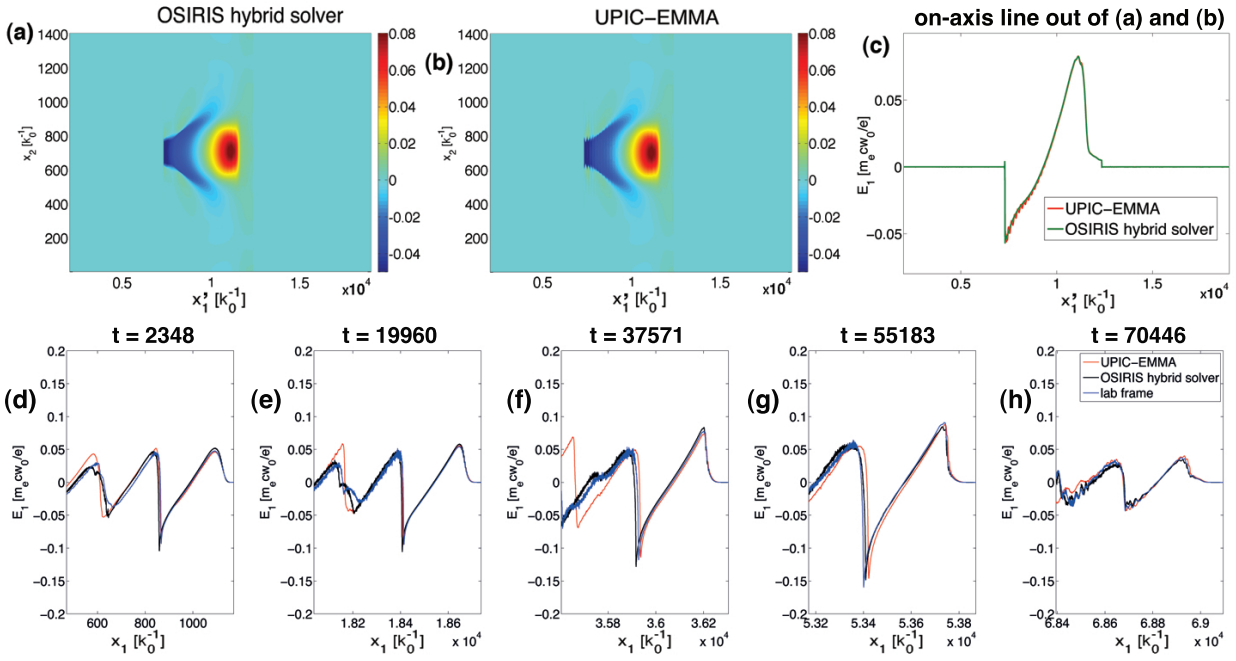

Figure 4.4: Comparison between OSIRIS lab frame, OSIRIS with the hybrid solver in the boosted frame and UPIC-EMMA in the boosted frame. In (a) and (b), 2D plots of E_1 for OSIRIS with the hybrid solver and UPIC-EMMA at $t' = 3955\omega_0^{-1}$ are shown in the boosted frame, where ω_0 is the laser frequency in the lab frame. In (c), line outs along the laser propagation direction of the same data are shown. In (d)–(h), line outs of the E_1 data transformed back to the lab frame are shown. The colored lines correspond to an OSIRIS lab frame simulation, an OSIRIS hybrid solver simulation in the Lorentz boosted frame, and UPIC-EMMA simulation in the Lorentz boosted frame.

the spatial derivatives in the $\hat{x}_1$ direction, while keeping the locality of the field solve and current deposit in the directions transverse to $\hat{1}$. For the current deposit, we start from the charge conserving deposit in OSIRIS and then correct it so that it still satisfies the continuity equation for the hybrid solver. Thus, Gauss's law remains rigorously satisfied at every time step if it is satisfied initially.

It is found from the NCI theory that such a hybrid solver has similar NCI properties in comparison to a full spectral solver that solves Maxwell's equations in multi-dimensional $\vec{k}$ space. As a result, the $(\mu, \nu_1) = (0, 0)$ NCI modes have a growth rate one order of magnitude smaller than the fastest growing $(\mu, \nu_1) = (0, \pm 1)$ NCI modes, and are highly localized. In addition, the growth rates of the $(\mu, \nu_1) = (0, 0)$ modes decrease as one reduces the simulation time step, and their locations in Fourier space also move farther away from the physics.

Compared with the spectral solver, the hybrid solver performs an FFT only along the drifting direction of the plasma. As a result, it eliminates the computation of FFTs in other direction(s). In addition, it can be readily adapted into fully operational FDTD codes without the need to modify various boundary conditions in the transverse directions. Very importantly, this idea can be readily applied to the quasi-3D algorithm in which the quantities are decomposed into azimuthal harmonics. In this algorithm FFTs cannot be used in the $\hat{r}$ direction. We demonstrate the feasibility of the hybrid Yee-FFT solver in 2D/3D Cartesian geometry, as well as in the quasi-3D geometry. Although we have not conducted a rigorous theoretical analysis for the NCI in the r-z or quasi-3D geometries, we find in simulations that the hybrid solver in quasi-3D geometry has very similar NCI properties to that in the 2D Cartesian geometry. This scheme requires a 1D FFT which cannot be parallelized at the moment.

We show that the strategy to eliminate NCI in the hybrid solver for 2D/3D Cartesian geometry, as well as quasi-3D geometry, is similar to that for the spectral solver. The fastest growing NCI modes can be eliminated by applying a low-pass filter in the current. The $(\mu, \nu_1) = (0, 0)$ NCI modes can be eliminated by reducing the time step which both reduces their growth rates and moves them away from the physical modes in Fourier space. These NCI modes can also be fully eliminated by slightly modifying the the EM dispersion relation

along k_1 direction at the location in Fourier space where the $(\mu, \nu_1) = (0, 0)$ modes reside. This approach is demonstrated in both Cartesian and quasi-3D geometry.

We showed that the new hybrid solver in OSIRIS can be used to conduct 2D LWFA simulations in a Lorentz boosted frame. With the low-pass filter applied to current and using a reduced time step, we observe no evidence of NCI affecting the physics in the simulation. Very good agreement is found between the results from OSIRIS with the hybrid solver, UPIC-EMMA simulations, as well as OSIRIS lab frame simulations with the standard Yee solver. This demonstrates the feasibility of using the hybrid solver to perform high fidelity simulations when plasma or beams drift relativistically across the grid.

CHAPTER 5

Customized FDTD Maxwell solver for NCI elimination

In this chapter, we take advantage of the previous progress described in previous chapters [46, 72], and develop a method to design a finite-difference-time-domain (FDTD) solver that has similar (yet different) NCI properties to the FFT-based solver described in Chapter 3 and 4 [46, 72]. Although it was based on the use of FFTs, when examined more carefully the previous progress showed that the key to essentially eliminate the NCI is to first isolate the range of unstable $\vec{k}$'s for what we refer to as the main NCI mode. This is accomplished by using a solver that has sufficiently small numerical errors in the spatial derivatives (and thus small dispersion errors for light waves) for moderate $|k_1|$. Even with perfect dispersion for light waves in vacuum, there will always be an intersection from the first spatial aliased beam modes at high $|k_1|$ that needs to be filtered out, and coupling between the EM mode and the main Langmuir mode at moderate $|k_1|$. We note that for this reason the use of the PSATD solver described in Ref. [41] does not appear to have advantages with respect to eliminating the NCI. As we have recently shown [46], when the main mode is isolated to a small range of $\vec{k}$s then a small modification to the dispersion for the range of unstable modes can remove the coupling between the EM (purely transverse in the lab frame) modes and Langmuir (purely longitudinal in the lab frame) modes.

Recognizing how the NCI is being controlled and eliminated by the use of an FFT along the plasma drifting direction ($\hat{1}$ direction) leads us to consider the possibility to design a customized and higher order finite difference operator for the spatial derivatives that provides sufficiently accurate dispersion for moderate $|\vec{k}|$. This finite difference operator for the spatial derivative can be implemented into a FDTD solver which is purely local and should thus scale well on massively parallel computers using domain decomposition. While this new

solver can eliminate unstable modes at moderate $|k_1|$, it cannot eliminate modes at high $|k_1|$ near the edge of the first Brillouin zone. In addition, such a solver will not conserve charge i.e., Gauss's law will not be satisfied. Nonetheless, both of these issues can be resolved by performing local FFTs for the current which do not use any global communication.

Similar to the hybrid Yee-FFT solver, to ensure the Gauss' Law is satisfied for the customized solver, we correct the component of the current in the $\hat{1}$ direction in k_1 space. This is done locally on each parallel partition along $\hat{1}$ and because the current is already in k_1 space, we can also use a low pass filter and eliminate the unstable high k_1 NCI modes. This filter can also be included into the current correction. We will show that overall this method is effective at eliminating the NCI while allowing good parallel scalability when domain decomposition is required along $\hat{1}$ direction.

We note that in Ref. [54] a PIC algorithm based on using a "local" FFT Maxwell solver was proposed. Their work was motivated for maintaining high parallel efficiency and was not focused on eliminating the NCI. They did show results from LWFA simulations in a boosted frame in a linear regime, however, no analysis for the NCI for FFT based algorithms was presented. We note that there are distinct differences between their approaches to ours. In our case, FFTs are performed only on the current. This is done to ensure that the continuity equation is satisfied, and also to filter the current for NCI elimination. Because the current from a single particle is not global this can lead to a current that satisfies the continuity equation at every location and it can eliminate the NCI. The EM fields remain in real space and are advanced using Faraday's and Ampere's law. So long as the solution for the current satisfies the continuity equation locally for the finite difference operators used in Ampere's Law, then Gauss's law will be maintained. If the longitudinal components of the fields are also solved using local FFTs in each partition (as is proposed in Ref. [54]), there will be errors in the longitudinal components of $\vec{E}$ and $\vec{B}$ due to the enforcement of periodicity.

The remainder of this chapter is organized as follows. In section 5.1 we first present a method to construct a customized FDTD Maxwell solver that has preferable NCI properties. The corresponding current correction and filtering strategies are discussed in section 5.2. We show that the use of local FFTs can provide a current that satisfies the continuity equation

for a customized solver. We then present sample simulations in section 5.3 showing that good accuracy and scalability can be obtained. Finally, in section 5.4 a summary is given.

5.1 Customized Maxwell solver

The Numerical Cerenkov Instability (NCI) occurs when a plasma drifts relativistically on a grid in a PIC code due to the unphysical coupling between the Langmuir modes (both main and aliasing) and electromagnetic (EM) modes [46]. Categorizing the NCI modes with their temporal aliasing mode number μ, and spatial aliasing mode number ν_1, it is found that usually the most violently growing NCI modes are those at $(\mu, \nu_1) = (0, \pm 1)$ (we call them first spatial aliasing NCI modes), and $(\mu, \nu_1) = (0, 0)$ (main NCI modes) [46]. The first spatial NCI modes usually reside near the edges of the fundamental Brillouin zone, making them relatively easy to eliminate with a sharp low-pass filter to the current. On the other hand, the main NCI modes are usually located within the inner half of the $\vec{k}$ modes, where modes of physical interest are located. It was shown for typical FDTD solvers that these modes were contained in a broad spectrum such that they cannot be eliminated through a low pass or mask filter. On the other hand, as shown in Chapter 3 and 4 [46, 72], for FFT-based solvers (and cell sizes $\Delta x_1 \leq \Delta x_2$), the main NCI modes are very localized in $\vec{k}$ space and they move to large $|\vec{k}|$ as the time step is reduced.

Therefore, as discussed in [46, 72], to eliminate the NCI modes in FFT-based solvers, the first step is to find a reduced time step which moves the main NCI modes away from the physical modes. After these modes are far enough from the physical modes, one can then apply a highly localized modulation to the EM dispersion relation in $\vec{k}$ space where the $(\mu, \nu_1) = (0, 0)$ modes reside in order to completely eliminate them. For a multi-dimensional spectral solver, the modification of the EM dispersion is accomplished by directly modifying the finite difference operator $[\vec{k}]$ in that localized area in $\vec{k}$ space [46]. For the hybrid Yee-FFT solver, only the operator $[k_1]$ is modified in the k_1 range where the main NCI modes are located [72]. The modification of the operator usually creates a bump in the dispersion curve in the range where the main NCI modes are located, which removes the coupling between

the EM modes and the Langmuir modes in that area, thereby eliminating the main NCI modes in that range completely.

When solving the Maxwell's equation using the FFT-based solvers, the differential operators in $\vec{k}$ space are explicitly used in the equations, therefore it is straightforward to modify the operators in $\vec{k}$ space. However, as mentioned in introduction, the use of a 1D FFT when there are many cells along that direction affects the scalability of the solver (or a multi-dimensional FFT solver when there are "many" more cells along one direction). Therefore, a question that naturally follows is whether it is possible to design an FDTD solver to imitate the characteristics of the EM dispersion curves of a FFT-based solver that make it possible to effectively eliminate the NCI. In the following sections, we describe a "recipe" for designing a finite difference derivative that when written in $\vec{k}$ space leads to the proper characteristics for the EM dispersion.

In [72] it is found that by replacing the finite difference spatial derivative in the direction of the plasma drift from a stencil that is second order accurate in cell size with a spectral solver (which is greater than N-th order accurate), one can restrict the $(\mu, \nu_1) = (0, 0)$ NCI modes to a highly localized area in the fundamental Brillouin zone [72]. Meanwhile, the spatial derivatives in the other direction(s) can be kept as second order accurate. Therefore, when we design an FDTD solver for the purpose of NCI elimination, it is natural to start with a "hybrid" FDTD solver that resembles the hybrid Yee-FFT solver. We use a higher order FDTD finite difference stencil [75, 76] in the direction of the drift while keeping them second order accurate in the other direction(s). Examination of the NCI growth rate where $[k_1]$ is replaced with the expression for a higher order stencil reveals that indeed the NCI is localized. In addition, we find that new $(\mu, \nu_1) = (0, 0)$ modes arise at large k_1 where the EM dispersion curve rolls over, i.e., the phase velocity drops. We then show how to modify the expression for $[k_1]$ for the higher order finite difference operator such that the EM dispersion curve has a slight bump at moderate $|k_1|$ in order to precisely avoid the coupling between the EM modes and main Langmuir modes for the main, $(\mu, \nu_1) = (0, 0)$, NCI modes of moderate $|k_1|$. To accomplish that, we expand the number of terms in the stencil [see Eq. (5.3)] to add extra degrees of freedom which can create the bump in the k_1 range where the main NCI

modes reside, as we will explain in the following sections. In addition the new $(\mu, \nu_1) = (0, 0)$ modes at high k_1 can be filtered out since they are outside the range of physically relevant modes.

5.1.1 NCI for high order finite difference solvers

Without loss of generality, we describe the method outlined above in the 2D Cartesian geometry. In a Maxwell solver, the electromagnetic fields $\vec{E}$ and $\vec{B}$ are advanced by solving Faraday's law and Ampere's law,

$$\vec{B}^{n+\frac{1}{2}} = \vec{B}^{n-\frac{1}{2}} - c\Delta t \nabla_p^+ \times \vec{E}^n \tag{5.1}$$

$$\vec{E}^{n+1} = \vec{E}^n + c\Delta t \nabla_p^- \times \vec{B}^{n+\frac{1}{2}} - 4\pi\Delta t \vec{J}^{n+\frac{1}{2}} \tag{5.2}$$

where the EM fields $\vec{E}$ and $\vec{B}$, and current $\vec{J}$ are defined on the staggered Yee grid [21], and $\nabla_p^\pm = \left(\partial_{p,x_1}^\pm, \partial_{2,x_2}^\pm\right)$ is the discrete finite difference operator for the staggered scheme. Note according to [72] the NCI can be eliminated if the operator $[k_1] = k_1$ is used along the plasma drifting direction. We now show that for a FDTD solver, a similar dispersion curve can be obtained by using high order finite difference operator in this direction. We apply a p-th order operator in the $\hat{1}$ direction and a standard second order Yee solver in the $\hat{2}$ direction. The p-th order operator is defined as

$$\begin{aligned}
\partial_{p,x_1}^+ f_{i_1,i_2} &= \frac{1}{\Delta x_1} \sum_{l=1}^{p/2} C_l^p \left(f_{i_1+l,i_2} - f_{i_1-l+1,i_2}\right) \\
\partial_{p,x_1}^- f_{i_1,i_2} &= \frac{1}{\Delta x_1} \sum_{l=1}^{p/2} C_l^p \left(f_{i_1+l-1,i_2} - f_{i_1-l,i_2}\right)
\end{aligned} \tag{5.3}$$

where f is an arbitrary quantity, i_1 and i_2 are the spatial grid indices, and the coefficients of the finite difference operator C_l^p can be expressed as [75, 76]:

$$C_l^p = \frac{(-1)^{l+1} 16^{1-\frac{p}{2}} (p-1)!^2}{(2l-1)^2 (\frac{p}{2}+l-1)! (\frac{p}{2}-l)! (\frac{p}{2}-1)!^2} \tag{5.4}$$

If we perform a Fourier transform to Eqs. (5.1) and (5.2) in both time and space, Maxwell's equations become

$$[\omega]\vec{B} = c[\vec{k}] \times \vec{E} \tag{5.5}$$

$$[\omega]\vec{E} = -c[\vec{k}] \times \vec{B} - 4\pi i \vec{J} \tag{5.6}$$

where the operators in operators in frequency and wavenumber space are

$$
\begin{aligned}
[\omega] &= \frac{\sin(\omega \Delta t/2)}{\Delta t/2} \\
[\vec{k}] &= \left(\sum_{l=1}^{p/2} C_l^p \frac{\sin[(2l-1)k_1 \Delta x_1/2]}{\Delta x_1/2}, \frac{\sin(k_2 \Delta x_2/2)}{\Delta x_2/2} \right)
\end{aligned}
\tag{5.7}
$$

where ω and $k_{1,2}$ are the frequency and wave numbers, and Δt and $\Delta x_{1,2}$ are the time step and grid sizes of the PIC system, respectively. Note that when the current vanishes, $\vec{J} = \vec{0}$, in Eq (5.6), we obtain the numerical dispersion relation for EM waves in vacuum, i.e.,

$$
[\omega]^2 = c^2 \left([k_1]_p^2 + [k_2]_2^2 \right) \tag{5.8}
$$

where $[k_1]_p$ and $[k_2]_2$ are the components of $[\vec{k}]$, and the order of accuracy is denoted by the subscripts outside the square brackets.

We plot the numerical dispersion relation ω v.s. k_1 (assuming $k_2 = 0$) in Fig. 5.1. We can see from Fig. 5.1 that, when the order p of $[k_1]$ increases, the dispersion curve is converging to (but never approaches) that of the spectral solver (black solid line). To quantify the locations and growth rates of the NCI modes for high order solvers, in Fig. 5.2, we plot the patterns of the $(\mu, \nu_1) = (0, 0)$ and $(0, 1)$ NCI modes over the (k_1, k_2) space in the fundamental Brillouin zone. The plot is generated by applying the p-th order finite difference operator in k_1 and second order finite difference operator in x_2 into the theoretical framework developed in Ref. [45, 46]. From Fig. 5.2 (a) we can see that the main NCI modes of a high order solver split into two parts: a highly localized part, i.e., a "dot", near $k_1/k_{g1} = 0.2$ (that has a lower growth rate), and another "strip" component that is very close to the edge of the fundamental Brillouin zone (that has a higher growth rate). To make both visible on the same scale we multiply the growth rate of the "dot" modes by ten. It is interesting to note that the highly localized "dot" part of the main NCI modes is located at almost the same place as for the hybrid Yee-FFT solver [shown in Fig. 5.2 (c)]. Meanwhile, the "strip" component has a growth rate on the same order of magnitude as the main NCI modes of the Yee solver, which are comparable to the $(\mu, \nu_1) = (0, 1)$ modes for either FFT or finite difference solvers [see Fig. 5.2 (b) and (d)]. This can be explained by the fact that in the low k_1 range the dispersion curve of the higher order solver almost overlaps that of the

hybrid solver, while for the high k_1 range the curve bends down, resulting in a similar NCI pattern to that of the Yee solver (which rolled over for lower k_1 values). In the meantime, the $(\mu, \nu_1) = (0, \pm 1)$ modes of the higher order solver reside very close to the edge of the fundamental Brillouin zone [shown in Fig. 5.2 (b)], which is similar to the case of the hybrid Yee-FFT solver [Fig. 5.2 (d)]. This enables us to readily eliminate these modes by applying a low-pass filter to the current in k_1-space.

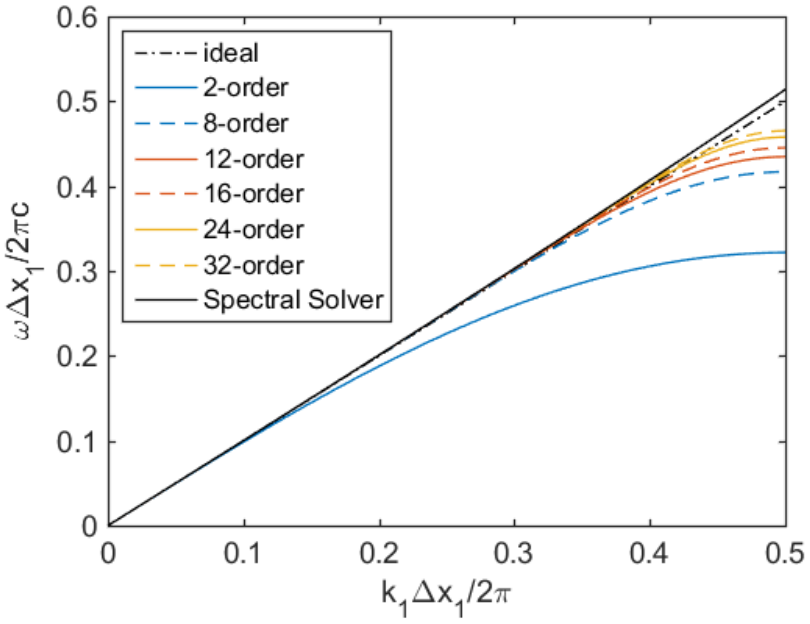

Figure 5.1: 1D numerical dispersion relations of different finite difference solver and spectral solver. $\Delta x_1 = 0.2k_p^{-1}$ and $\Delta t = 0.05\omega_p^{-1}$, where n_p is the reference plasma density, and $\omega_p^2 = 4\pi q^2 n_p/m_e$, $k_p = \omega_p$ (c is normalized to 1), are used to generate the plots.

Just as was the case for the hybrid Yee-FFT solver, the location of the "dot" part of the main NCI modes also changes for the higher order solver as one reduces the time step. In Fig. 5.3, we scan the location of the "dot" part of the main NCI modes with different time steps for various solvers. We can see that the main NCI modes at moderate $|\vec{k}|$ move towards higher $|\vec{k}|$ for both the hybrid Yee-FFT solver, 16th order solver, and 24th order solver. Therefore, it is possible to apply the strategies used for the hybrid Yee-FFT solver to effectively eliminate the NCI for the hybrid higher order-Yee solver. In addition, as $\Delta t/\Delta x_1$ decreases, so do the growth rates for the "dot" part of the main NCI modes.

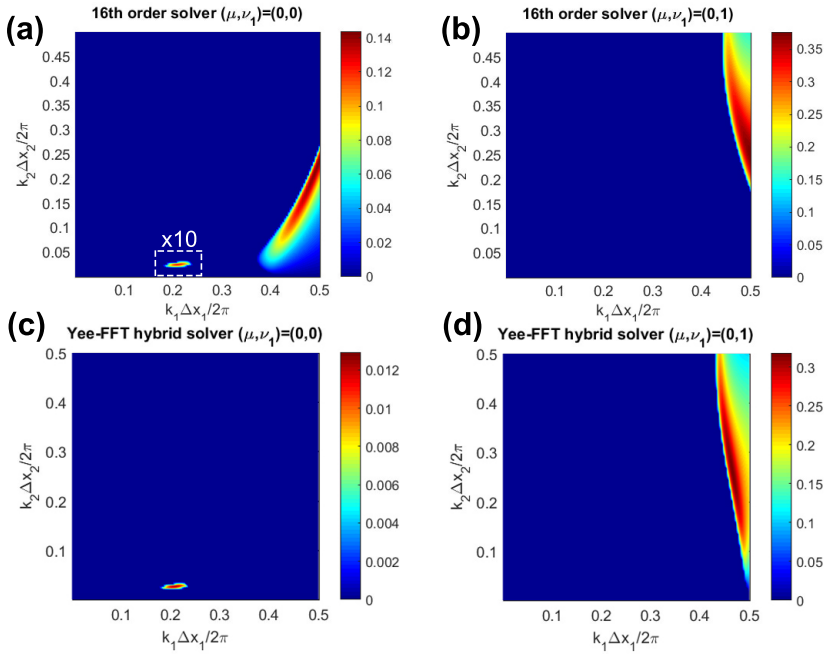

Figure 5.2: The NCI patterns of 16th order solver and hybrid Yee-FFT solver. (a) and (b) show the $(\mu, \nu_1) = (0, 0)$ and $(\mu, \nu_1) = (0, 1)$ NCI mode of the high order (16th order) solver. The values in the dashed line box in (a) is multiplied by 10 for better visualization. (c) and (d) show the $(\mu, \nu_1) = (0, 0)$ and $(\mu, \nu_1) = (0, 1)$ NCI mode of the hybrid solver. $\Delta x_{1,2} = 0.2k_p^{-1}$, $\Delta t = 0.05\omega_p^{-1}$ and $n = 30n_p$, where n_p is the reference plasma density, and $\omega_p^2 = 4\pi q^2 n_p/m_e$, $k_p = \omega_p$ (c is normalized to 1), are used to generate the plots.

For given simulation parameters, we first calculate the locations of the main NCI modes for the 16th order solver (16th order in $\hat{1}$ direction, and 2nd order in other directions). If they are too close to the physical modes, we reduce the time step to move them away from the center towards the edge of the fundamental Brillouin zone. In the next section, we describe how to modify the higher order stencil such that its k_1 v.s. $[k_1]_p$ curve has a bump to eliminate the "dot" part of the main NCI modes.

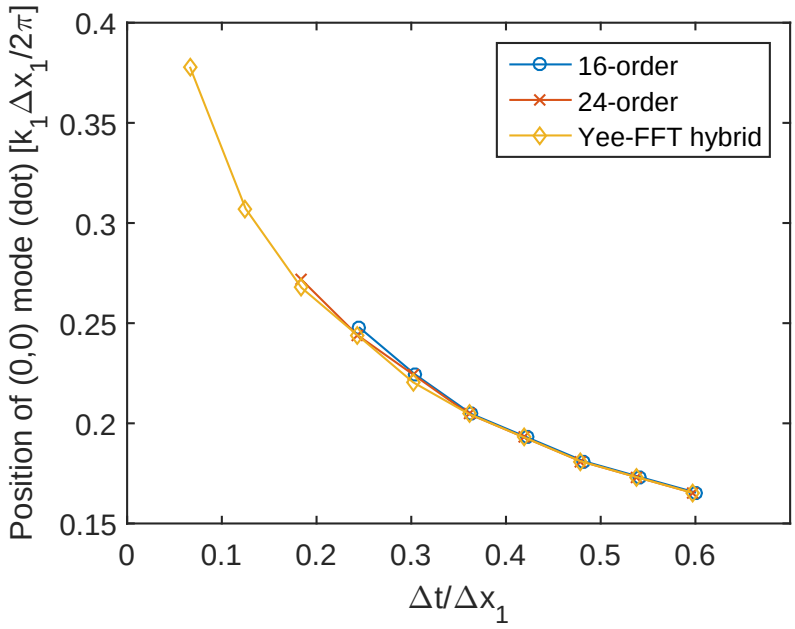

Figure 5.3: The position of the splitted "dot" in $(\mu, \nu_1) = (0, 0)$ NCI mode at different time step for 16th order, 24th order solver and Yee-FFT hybrid solver. We scan the position using $\Delta x_{1,2} = k_0^{-1}$ and $n_p = n_0$, from $\Delta t = 0.1\Delta t_{\text{CFL}}$ to $\Delta t = 0.9\Delta t_{\text{CFL}}$.

5.1.2 Customization of $[k]_1$

In this subsection, we explain how we customize a higher order finite difference first derivative that also has a slight modification near the location of the NCI modes in wave number space. For the FFT-based solvers described in [46, 72] this modification to the EM dispersion relation in the k_1 range where the main NCI modes are located can be easily implemented. Specifically, this is accomplished by changing the definition of k_1 inside the field solver

to $[k_1](k_1)$ in the range $k_1 \in [k_{1l}, k_{1u}]$, where the $(\mu, \nu_1) = (0, 0)$ NCI modes reside to $[k_1] = k_1 + \Delta k_{\text{mod}}(k_1)$ where $\Delta k_{\text{mod}}(k_1)$ is a small localized perturbation(see Fig. 3(a) in [72]). It usually takes the form of

$$\Delta k_{\text{mod}}(k_1) = \begin{cases} \Delta k_{\text{mod,max}} \sin\left(\pi \frac{k_1 - k_{1l}}{k_{1u} - k_{1l}}\right)^2, & k_{1l} \le k_1 \le k_{1u} \\ 0, & \text{otherwise} \end{cases} \tag{5.9}$$

where k_{1l}, k_{1u} are the lower and upper bounds of the region that is modified, and $\Delta k_{\text{mod,max}}$ is the maximum of Δk_{mod}. In an FFT-based solver this modification is straightforward to implement, while in a high order FDTD solver, one has to find a stencil that has both higher order accuracy for the derivative over a wide range of wave number space as well as a modification in a local region of wave number space. For a regular pth order solver (where p is an even number), there are $p/2$ coefficients, C_l^p, for the stencil and the numerical dispersion relation is uniquely determined. It naturally follows that if we want to customize the dispersion relation based on the high order solver, we will need to add more degrees of freedom, i.e., more coefficients, into the operator. This means the stencil will be broader. The coefficients will still need to be constrained so that the operator has higher-order accuracy, while at the same time it has the desired modification in a local region of k_1 space.

We denote the high order solver as $\nabla_{p*}^{\pm} = \left(\partial_{p*,x_1}^{\pm}, \partial_{2,x_2}^{\pm},\right)$. The first component has the form

$$\partial_{p*,x_1}^{+} f_{i_1,i_2} = \frac{1}{\Delta x_1} \sum_{l=1}^{M} \tilde{C}_l^p \left(f_{i_1+l,i_2} - f_{i_1-l+1,i_2}\right) \tag{5.10}$$

$$\partial_{p*,x_1}^{-} f_{i_1,i_2} = \frac{1}{\Delta x_1} \sum_{l=1}^{M} \tilde{C}_l^p \left(f_{i_1+l-1,i_2} - f_{i_1-l,i_2}\right) \tag{5.11}$$

while the second component is still the standard second order accurate operator. The modified solver has M coefficients, $\tilde{C}_l^p$, where $M > p/2$. The corresponding finite difference operator in k-space becomes

$$[k_1]_{p*} = \sum_{l=1}^{M} \tilde{C}_l^p \frac{\sin[(2l-1)k_1 \Delta x_1/2]}{\Delta x_1/2} \tag{5.12}$$

In order to construct the "bump" in the dispersion curve for the proposed solver, we need to find an "optimized" set of $\tilde{C}_l^p$ such that $[k_1]_{p*}$ will best approximate the modified $[k_1] =$

$[k_1]_p+\Delta k$ described in Eq. (5.9). For the purpose of simplifying the notation, in the following we normalize $[k_1]_{p*}$, $[k_1]_p$, $\Delta k_{\rm mod}$, k_{1l}, k_{1u} and k_1, with $k_{g1} = 2\pi/\Delta x_1$. In the spirit of the least square approximation, we construct a function F_1

$$F_1 = \int_0^{1/2} \left([k_1]_{p*} - [k_1]_p - \Delta k_{\rm mod}\right)^2 dk_1 \tag{5.13}$$

which we will minimize to find an optimum set of $\tilde{C}_j^p$. We note that a weight function $w(k_1)$ can be included to minimize errors in targeted region of k_1 space, and this is an area for future work. In addition, the high order solver should also meet the requirement of pth order accuracy and satisfy the condition $\partial^{\pm}_{p*,x_1} \to \partial_{x_1}$ as $\Delta x_1 \to 0$. Therefore, the coefficients should be subject to the linear equations below

$$\mathcal{M}\tilde{\vec{C}}^p = \hat{\vec{e}}_1 \tag{5.14}$$

where $\tilde{\vec{C}}^p = (\tilde{C}_1^p, ..., \tilde{C}_M^p)^T$, $\hat{\vec{e}}_1 = (1, 0, ..., 0)^T$ and the elements of the matrix $\mathcal{M}$ are $\mathcal{M}_{ij} = (2j-1)^{2i-1}/(2i-1)!$ $(i = 1, ..., p/2)$ and $(j = 1, ..., M)$. This is a constrained least-square minimization problem so we can use the Lagrange multipliers to solve it.

The Lagrangian is defined by $\mathcal{L} = F_1 + F_2$ where $F_2 = \vec{\lambda}^T(\mathcal{M}\tilde{\vec{C}}^p - \hat{\vec{e}}_1)$ and $\vec{\lambda} = (\lambda_1, ..., \lambda_{p/2})^T$. Solving the constrained least-square minimization problem is equivalent to solving,

$$\partial\mathcal{L}/\partial\tilde{C}_j^p = 0 \ (j = 1, ..., M) \text{ and } \partial\mathcal{L}/\partial\lambda_i = 0 \ (i = 1, ..., p/2) \tag{5.15}$$

It can be straightforward to show that this results in the following set of equations,

$$\frac{\partial F_1}{\partial \tilde{C}_j^p} = \begin{cases} \frac{1}{2\pi^2}(\tilde{C}_j^p - C_j^p - A_j), & 1 \le j \le \frac{p}{2} \\ \frac{1}{2\pi^2}(\tilde{C}_j^p - A_j), & \frac{p}{2}+1 \le j \le M \end{cases} \tag{5.16}$$

$$\frac{\partial F_2}{\partial \tilde{C}_j^p} = \sum_i \lambda_i \mathcal{M}_{ij} \tag{5.17}$$

$$\frac{\partial F_2}{\partial \lambda_j} = \sum_i \mathcal{M}_{ji} A_i - e_j \tag{5.18}$$

where

$$A_j = \frac{8\Delta k_{\rm mod,max}(\cos[(2j-1)\pi k_{1u}] - \cos[(2j-1)\pi k_{1l}])}{(2j-1)[(2j-1)^2(k_{1u}-k_{1l})^2 - 4]} \tag{5.19}$$

Combing Eqs. (5.16)-(5.18), we can reformat Eq. (5.15) into a matrix equation

$$\begin{pmatrix} \frac{1}{2\pi^2}\vec{I} & \mathcal{M}^T \\ \mathcal{M} & \vec{0} \end{pmatrix} \begin{pmatrix} \tilde{\vec{C}}^p \\ \vec{\lambda} \end{pmatrix} = \begin{pmatrix} \frac{1}{2\pi^2}(\vec{A}+\vec{C}^p) \\ \hat{\vec{e}}_1 \end{pmatrix} \tag{5.20}$$

where $\vec{I}$ is the $M \times M$ unit matrix and $\vec{C}^p = (C^p_1, ..., C^p_{p/2}, 0, ..., 0)^T$. For given "bump" parameters $\Delta k_{\text{mod,max}}$, k_{1l} and k_{1u}, Eq. (5.20) can be solved numerically. Henceforth, in this chapter we use $M = p$.

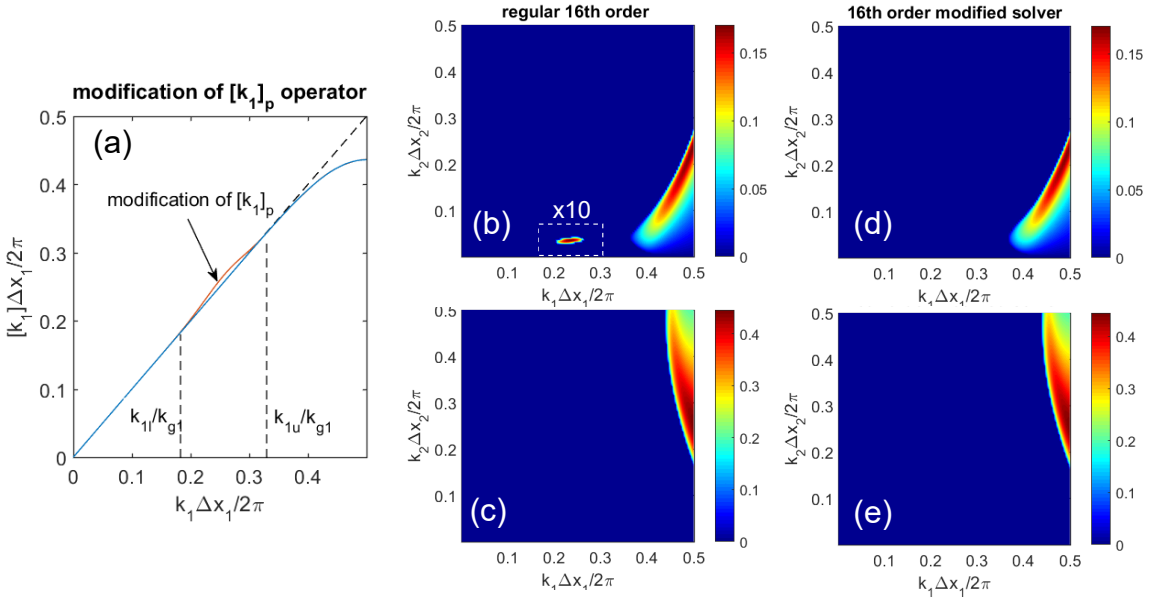

Figure 5.4: In (a) the perturbation (red line) introduced by the modified 16th order solver is shown. The blue line denotes the $[k_1]$-k_1 relation of the regular 16th order solver. The lower and upper limits of perturbation are $k_{1l}/k_{g1} = 0.18$ and $k_{1u}/k_{g1} = 0.33$. The perturbation magnitude $\Delta k_{\text{mod,max}} = 0.01$. (b) and (c) are $(\mu, \nu_1) = (0, 0)$ and $(\mu, \nu_1) = (0, 1)$ NCI mode patterns of the regular 16th order solver respectively. The values in the dashed line box in (b) are multiplied by 10 to make them more visible. (d) and (e) are the patterns of the modified 16th order solver. When generating these plots we use $\Delta x_{1,2} = 0.2k_p^{-1}$, $\Delta t = 0.05\omega_p^{-1}$ and $n = 50n_p$, where n_p is the reference plasma density, and $\omega_p^2 = 4\pi q^2 n_p/m_e$, $k_p = \omega_p$ (c is normalized to 1).

In Fig. 5.4 we show the comparison of the $[k_1]$ operator and the NCI mode patterns between the regular and customized high order solver. We use a 16th order solver as an

example. In Fig. 5.4(a), we show that a perturbation (red line) can be introduced to the $[k_1]_{16}$ operator within the bump region, while the operators $[k_1]_{16}$ and $[k_1]_{16*}$ are almost overlapped outside the region. Figs. 5.4 (b) and (c) show the NCI mode patterns for the regular 16th order solver. We can see the dot of $(\mu, \nu_1) = (0, 0)$ modes presented in the middle of the fundamental Brillouin zone, for which we aim to eliminate through the use of the modified solver. In Fig. 5.4(d), the dot has been eliminated and $(\mu, \nu_1) = (0, 1)$ modes and the remaining of $(\mu, \nu_1) = (0, 0)$ modes are almost identical to what was seen for the regular high order solver.

5.1.3 Courant condition

The derivation of the Courant condition for the proposed high order solver is straightforward. From the numerical dispersion relation

$$[\omega]^2 = c^2 \left([k_1]_{p*}^2 + [k_2]_2^2 \right), \tag{5.21}$$

it can be shown that to keep ω a real number, the corresponding constraint on the time step that must be satisfied is

$$\frac{\Delta t}{2} \sqrt{\left(\sum_{l=1}^{M} \tilde{C}_l^p \frac{\sin[(2l-1)k_1 \Delta x_1/2]}{\Delta x_1/2} \right)^2 + \left(\frac{\sin(k_2 \Delta x_2/2)}{\Delta x_2/2} \right)^2} \leq 1. \tag{5.22}$$

Note that $|k_1| \leq \pi/\Delta x_1$ and $|k_2| \leq \pi/\Delta x_2$, therefore the Courant condition of the proposed high order solver

$$\Delta t \leq 1/\sqrt{\frac{\left(\sum_{l=1}^{M} \tilde{C}_l^p \right)^2}{\Delta x_1^2} + \frac{1}{\Delta x_2^2}} \tag{5.23}$$

For the standard high order solver, given the cell sizes, the Courant limit only depends on cell sizes and the order solver's accuracy. For instance, the Courant limit of a 16th order solver (16th order in x_1, while second order in x_2) is $\Delta t_{CL} = 0.6575 \Delta x_1$ with $\Delta x_1 = \Delta x_2$. As for the customized solver, although the specified solver coefficients $\tilde{C}_l^p$ depend on the modification of numerical dispersion we introduce, the Courant limit on the time step varies little as we alter the "bumps" in the numerical dispersion curve. Taking the 16th order customized solver with 16 coefficients and $\Delta x_1 = \Delta x_2$ as an example, the Courant condition

reduces to $\Delta t_{CL} = 0.6550\Delta x_1$ for $k_l = 0.1$, $k_u = 0.3$, $\Delta k_{\mathrm{mod}} = 0.01$, and $\Delta t_{CL} = 0.6562\Delta x_1$ for $k_l = 0.15, k_u = 0.3, \Delta k_{\mathrm{mod}} = 0.005$. As we can see, the Courant condition changes little when switching from the standard high order solver to the customized solver.

5.1.4 Cartesian 3D and quasi-3D scenarios

As can be seen from previous sections, this FDTD solver only modifies the finite difference operator in the plasma drifting direction. As a result, although not presented in this chapter, the method described above can be extended to Cartesian 3D and quasi-3D geometry [51, 55] in a straightforward way.

5.2 Charge conservation and parallelization of the solver

Similar to the case of the hybrid Yee-FFT solver, when the $[k_1]$ of the solver is different from the second order accurate $[k_1]_2$, one needs to apply a correction to the current in order to satisfy Gauss' Law. This is due to the fact that in a typical FDTD EM-PIC code, the EM fields are advanced by the Faraday's law and Ampere's law, while the Gauss' Law is satisfied by applying a charge conserving current deposit [23]. The charge conserving current deposition is second order accurate in all directions, which matches exactly to the standard Yee solver. As a result, when the finite difference operator for the derivative along a particular direction changes in a solver, Gauss' Law is no longer satisfied if the current is not corrected correspondingly. Currently no rigorous charge conserving current deposit exists for higher order solvers.

More specifically, the charge conserving current deposition ensures the second-order-accurate finite difference representation of the continuity equation,

$$\overline{\frac{\partial}{\partial t}}\rho^n + \nabla_2^- \cdot \vec{J}^{n+\frac{1}{2}} = 0 \tag{5.24}$$

where $\overline{\frac{\partial}{\partial t}} G^n = \frac{G^{n+1} - G^n}{\Delta t}$ for an arbitrary scalar quantity G^n. For the Yee solver, Gauss's law is rigorously satisfied every time step if it is satisfied at the beginning. However, when combining the high order solver and the second-order-accurate current deposition scheme,

we need to apply a correction to the current in the drifting direction in order that Gauss's law remains satisfied throughout the whole simulation. After the current has been calculated locally, we then "correct" them by performing an FFT along the x_1 direction,

$$\tilde{J}_1^{n+\frac{1}{2}} = \frac{[k_1]_2}{[k_1]_{p*}} J_1^{n+\frac{1}{2}} \tag{5.25}$$

where $\tilde{J}_1$ is the corrected current, in a similar manner to what was employed for the hybrid Yee-FFT solver. Performing Fourier transform in the x_1 direction and applying the correction scheme in Eq. (5.25), we guarantee

$$\overline{\frac{\partial}{\partial t}} \rho^n(k_1, x_2) + i[k_1]_{p*} \tilde{J}_1^{n+\frac{1}{2}}(k_1, x_2) + \partial_{2,x_2}^- J_2^{n+\frac{1}{2}}(k_1, x_2) = 0 \tag{5.26}$$

Combining Eq. (5.26) with Ampere's law, Eq. (5.2), (replacing ∇_p^- with ∇_{p*}^- to be consistent with the modified high order solver), we obtain

$$\overline{\frac{\partial}{\partial t}} \left(-4\pi\rho^n + i[k_1]_{p*} E_1^n + \partial_{2,x_2}^- E_2^n \right) = 0 \tag{5.27}$$

We carry out inverse Fourier transform to retrieve the equation in real space,

$$\overline{\frac{\partial}{\partial t}} \left(-4\pi\rho^n + \nabla_{p*}^- \cdot \vec{E}^n \right) = 0 \tag{5.28}$$

which indicates that Gauss's law is satisfied if it is satisfied initially.

It is important to note that one effect that arises from the current correction is that the current from one particle extends to more cells. Therefore, an originally localized current distribution would be spread out over more cells after we correct the current in the k-space and transform back to real space. This results from the use of a less local operator for the derivative. Nevertheless, the current still rigorously satisfies the continuity equation for the desired particle shape. We point out that this effect can be neglected and the current correction is still a nearly error-free scheme in some sense, as will be discussed below.

Assuming we have a point current initially in the real space, located at the grid index $i_1 = 0$, *i.e.* $J_{1,i_1} = \delta(i_1)$. After performing the discrete Fourier transform, the current in the k-space becomes unity for all k_1, *i.e.* $J_{1,\kappa_1} = 1$, where $\kappa_1 = -N/2, ..., N/2 - 1$ is the mode

number, and $k_1 = 2\pi\kappa_1/N\Delta x_1$ with N the number of cells in x_1 direction. The corrected current $\tilde{J}_{1,\kappa_1}$ is therefore the correction factor

$$\left(\frac{[k_1]_2}{[k_1]_{p*}}\right)_{\kappa_1} = \frac{\sin\left(\frac{\pi}{N}\kappa_1\right)}{\sum_l \tilde{C}_l^p \sin\left[\frac{(2l-1)\pi}{N}\kappa_1\right]} \tag{5.29}$$

It can be asserted that the inverse discrete Fourier transform (IDFT) of the correction factor extends over all space. This can be easily proved by contradiction. Suppose the IDFT of the corrected point current was confined to a region of space, e.g., it is non zero only in the interval $[-W, W]$ and the expanded current distribution in the real space is symmetric, *i.e.* $J_{1,i_1} = J_{1,-i_1}$. The corresponding Fourier coefficients after a DFT would therefore be a finite summation of cosine functions, *i.e.* $J_{1,\kappa_1} = J_{1,i_1=0} + 2\sum_{i_1=1}^{W} J_{1,i_1}\cos(2\pi i_1\kappa_1/N)$. By inspection, it is clear that, Eq. (5.29) cannot be rewritten in such way. Therefore, we have proven that the IDFT of the current correction for a delta function cannot be localized in space, and therefore must exist within the entire space over which the FFT is performed.

However, as we show next the values of the corrected current for an initial delta function (at the origin) fall off rapidly for cells away from the origin. In fact, the values fall below double precision accuracy effectively making the corrected current effectively only non-zero in a finite region. We illustrate this through 1D numerical calculations. We initialize the current as a Dirac-delta distribution in the center of the grid, *i.e.* $J_1 = \delta(x_0)$. Then we perform a 1D FFT to the current, then use the correction according to Eq. (5.25) for solvers with different orders in the k-space, and finally perform an IFFT to transform the current back into real space. Fig. 5.5 shows the spatial distribution of the corrected current. It can be seen that the extent of the current is widened by the correction process and the width increase as the order of the solver is increased. Nonetheless, the expansion remains in a localized region in real space and beyond this region the amplitude is on the order of 10^{-16}, which corresponds to double precision roundoff. The fact that the current is localized indicates that the current correction can be done on a local domain so long as a sufficient number of guard cells are used. This permits of using domain decomposition along the $\hat{1}$ direction and the use of a local FFT on each domain. For example, we have decreased the size of the domain from 256 grids in Fig. 5.5 (a) to 128 grids in Fig. 5.5 (b). The

current distribution, the width expansion, and noise level are almost the same. Therefore, we can see that the size of a parallel partition makes little impact on the simulation results. Although the current expansion is mathematically infinite as mentioned before, in the sense of considering the precision of the numerical algorithm, the current expansion can be viewed as localized. When a sufficient number of guard cells for each simulation partition is used, this current correction scheme will be nearly free of any error brought by the current expansion.

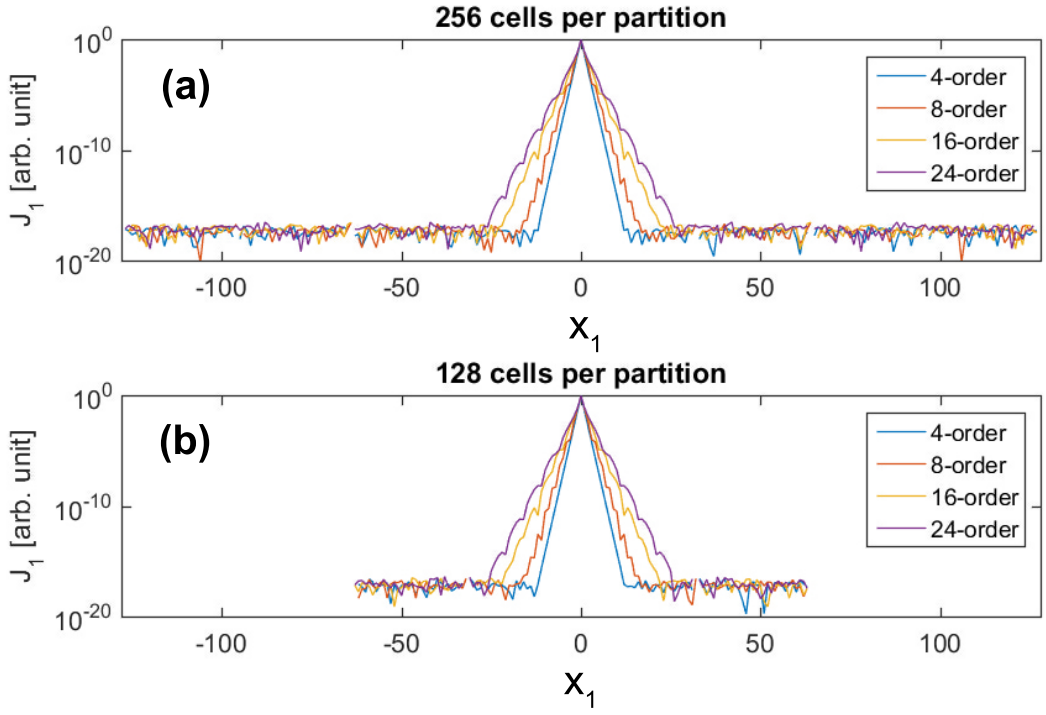

Figure 5.5: Effect of the current expansion tested by point current. Numerical calculations are carried out on 256 cells (a) and 128 cells (b), to model the cases using different partition sizes. Current with Dirac-delta distribution is initialized and the current corrections of different orders of solver are applied in the k-space. (a) and (b) show the current distributions in real space with different correction schemes. We set $\Delta x_1 = 1$ for the calculations.

5.3 Sample simulations

In this section we present sample simulations using the customized solver and its corresponding NCI elimination schemes. In these simulations, we use a low-pass filtering for the current

as described in Eq. (4.31).

5.3.1 Drifting plasma

In this subsection we demonstrate the capability of the customized FDTD Maxwell solver to eliminate the NCI using 2D Cartesian OSIRIS simulations of drifting plasmas. We fill the simulation box with a plasma drifting relativistically at $\gamma = 20$ along $\hat{1}$ direction. The plasma has a very small but finite temperature to seed the NCI. Periodic boundary conditions are used for both the $\hat{1}$ and $\hat{2}$ directions. We performed simulations using both the 16th order solver, the customized solver, and the hybrid Yee-FFT solver, with and without the low-pass filters. Other simulation parameters are presented in Table 5.1, and the corresponding coefficients for the customized solver are listed in Table 5.2.

Parameters	**Values**
grid size $(\Delta x_1, \Delta x_2)$	$(0.5k_p^{-1}, 0.5k_p^{-1})$
time step Δt	$0.25\Delta x_1$
number of grid	512×512
particle shape	quadratic, cubic
electron drifting momentum p_{10}	19.975 $m_e c$
plasma density	2.0 n_p

Table 5.1: Simulation parameters for the 2D drifting plasma simulation. n_p is the plasma density, and $\omega_p^2 = 4\pi q^2 n_p/m_e$, $k_p = \omega_p$ (c is normalized to 1).

We can see from Fig. 5.6 that, by comparing the red line (16th order solver without any filters) and the orange line (16th order solver plus low-pass filter), that applying the low-pass filter to a 16th order solver significantly reduces the growth of the E_2 energy. This is because the fastest growing $(\mu, \nu_1) = (0, \pm 1)$ NCI modes are eliminated by the low-pass filter. Besides using the low-pass filter, when we add the bump to the 16th order solver (thus making it a customized solver), the growth rate is further reduced since the main NCI modes are completely eliminated (see the blue line). Even higher order NCI modes [46] are

Coefficients	Values	Coefficients	Values
$\tilde{C}_1^{16}$	1.237042976225048	$\tilde{C}_2^{16}$	-0.102548201854464
$\tilde{C}_3^{16}$	0.022015354460742	$\tilde{C}_4^{16}$	-0.009258452621442
$\tilde{C}_5^{16}$	0.000410036656959	$\tilde{C}_6^{16}$	0.002572239519500
$\tilde{C}_7^{16}$	0.001482836071727	$\tilde{C}_8^{16}$	-0.001392055950412
$\tilde{C}_9^{16}$	-0.001472515326959	$\tilde{C}_{10}^{16}$	0.000478783514362
$\tilde{C}_{11}^{16}$	0.001200462462019	$\tilde{C}_{12}^{16}$	-0.000187062256742
$\tilde{C}_{13}^{16}$	-0.001059471474041	$\tilde{C}_{14}^{16}$	0.000873314953435
$\tilde{C}_{15}^{16}$	-0.000281855449164	$\tilde{C}_{16}^{16}$	0.000034281167855

Table 5.2: Coefficients $\tilde{C}_i^{16}$ in Eq. (5.12) for the customized solver based on the 16th order solver, for the single plasma drift simulation discussed in section 5.3.1.

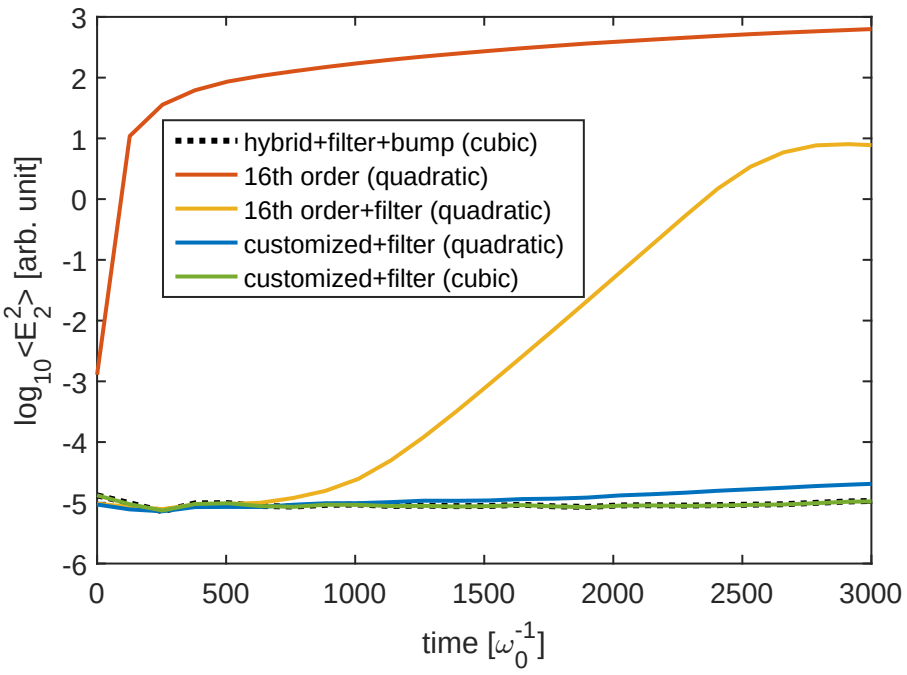

Figure 5.6: Evolutions of the E_2 energies in the numerical systems for various setups in drifting plasma 2D Cartesian PIC simulations, as discussed in section 5.3.1.

attributed to the slight growth in energy for this case (blue line in Fig. 5.6), and when the cubic particle shape is applied, the corresponding energy growth is effectively suppressed (green line in Fig. 5.6). We can see that when the low-pass filter, bump, and higher order particles are applied to the cases of both customized solver (green line) and hybrid Yee-FFT solver (black dotted line), both the hybrid Yee-FFT solver and customized solver schemes effectively eliminate the NCI.

5.3.2 Relativistic shock

We next present an example where two plasmas are collided against each other which is relevant for relativistic shock simulations. The two plasmas drift towards each other with a Lorentz factor of $\gamma = 20.0$. The simulation has a box size of 131072×2048 cells, for which the number of cells in $\hat{1}$ direction is much larger than that of the $\hat{2}$ direction. Since the plasmas are drifting in the $\hat{1}$ direction, the total number of cores that can be used in such a simulation would be significantly limited if we use FFT-based solvers, which requires one partition along the $\hat{1}$ direction. With the customized solver and corresponding elimination scheme, we are able to partition in the $\hat{1}$ direction. We used a 2D domain decomposition with 256×16 partitions along the $\hat{1}$ and $\hat{2}$ directions respectively. Other simulation parameters are listed in Table 5.3, and the corresponding coefficients for the customized solver are listed in Table 5.4.

In Fig. 5.7 we plot the 2D color isosurface plots of the ion density, and line outs of the x_2 averaged ion density for the Yee solver, and customized solver. For the Yee solver case, we used the optimal time step of $\Delta t = 0.5\Delta x_1$ at which the NCI is minimized [40, 45], plus a 5-pass current smoothing and compensation for the current, and the EM fields are also filtered every two time steps. For the case with the customized solver, we used the same NCI elimination scheme as is used for the single drifting plasma case, but with slightly different coefficients (different density and time step) for the solver. We can see from Fig. 5.7 that there are noticeable differences. For example, the transverse size of the filaments in the density are larger for the Yee case than for the customized case. From our NCI

theory we know that the Yee solver with the optimized time step does not eliminate the main $(\mu, \nu_1) = (0, 0)$ modes. The growth rate for these modes is reduced, but they are not localized in space; instead they reside within the range of physics. It is not obvious when, and how these modes are altering the physics. On the other hand, the customized solver together with our filters completely removes the $(\mu, \nu_1) = (0, \pm 1)$ and $(\mu, \nu_1) = (0, 0)$ modes; and the use of the higher order particle shapes reduces the growth rate for the next highest growing modes.

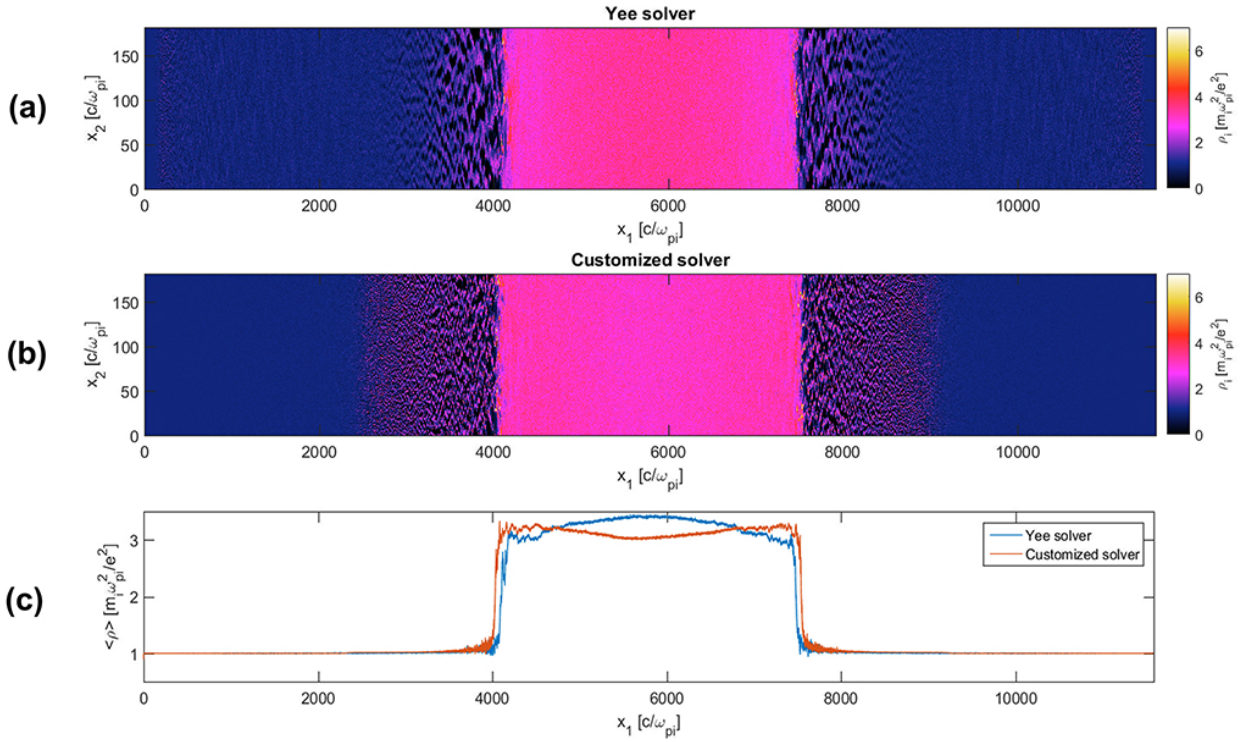

Figure 5.7: The ion densities and their line outs for a relativistic shock simulation, as discussed in section 5.3.2. The corresponding simulation parameters are listed in Table 5.3.

5.3.3 LWFA boosted frame simulation

In this subsection, we present 3D Cartesian LWFA boosted frame PIC simulations using the customized FDTD Maxwell solver in OSIRIS. For comparison, we also carried out simulations using the hybrid Yee-FFT solver and customized FDTD Maxwell solver respectively. The parameters correspond to the lab frame simulation discussed in [16] in which 1.3 GeV

Parameters	**Values**
Plasma	
density n_{pe}, n_{pi}	n_0
initial Lorentz factor γ_0	20.0
initial thermal velocity $v_{\mathrm{th},e,i}$	$8.7 \times 10^{-5} c$
mass ratio m_i/m_e	32
Simulation using customized high order solver	
cell size $\Delta x_{1,2}$	$0.5 k_0^{-1}$
time step $\Delta t/\Delta x_1$	0.2
number of cells	$2^{17} \times 2^{11}$
particle shape	cubic
particle per cell	$(1, 2)$
$[k_1]$ modification $(k_{1l}, k_{1u}, \Delta k_{\mathrm{mod,max}})$	$(0.1, 0.35, 0.01)$
low-pass filter (f_l, f_u)	(0.275,0.3)
Simulation using standard Yee solver	
cell size $\Delta x_{1,2}$	$0.5 k_0^{-1}$
time step $\Delta t/\Delta x_1$	0.5
number of cells	$2^{17} \times 2^{11}$
particle shape	cubic
particle per cell	$(1, 2)$

Table 5.3: Parameters for 2D relativistic collisionless plasma simulations in lab frame using the modified high order solver and Yee solver. The plasma density n_0 and corresponding wave number k_0 are used to normalize the simulation parameters. The parameters of $[k_1]$ modification are normalized to $k_{g1} \equiv 2\pi/\Delta x_1$.

Coefficients	**Values**	**Coefficients**	**Values**
$\tilde{C}_1^{16}$	1.243205632406442	$\tilde{C}_2^{16}$	-0.096527073844747
$\tilde{C}_3^{16}$	0.017018941335700	$\tilde{C}_4^{16}$	-0.013839950216042
$\tilde{C}_5^{16}$	0.003588768352855	$\tilde{C}_6^{16}$	0.005153133591937
$\tilde{C}_7^{16}$	0.000007068893273	$\tilde{C}_8^{16}$	-0.002317133408538
$\tilde{C}_9^{16}$	-0.001166192174494	$\tilde{C}_{10}^{16}$	0.000552266782136
$\tilde{C}_{11}^{16}$	0.001508596910066	$\tilde{C}_{12}^{16}$	-0.000134050410326
$\tilde{C}_{13}^{16}$	-0.001599956501178	$\tilde{C}_{14}^{16}$	0.001305552125425
$\tilde{C}_{15}^{16}$	-0.000423469804615	$\tilde{C}_{16}^{16}$	0.000051829248350

Table 5.4: Coefficients $\tilde{C}_i^{16}$ in Eq. (5.12) for the customized solver based on the 16th order solver, for the relativistic shock simulations, and LWFA simulations in the Lorentz boosted frame, as discussed in section 5.3.2 and section 5.3.3.

electrons are generated. We have also listed the parameters in Table 5.5. Note that although the simulation parameters in this scenario is different from those of the relativistic shock simulations discussed in section 5.3.2, the locations of the main NCI modes for a 16th order solver under these two sets of parameters are very close to each other. Therefore we used the same coefficients for the customized solver as in section 5.3.2, as listed in Table 5.4.

In Fig. 5.8(a) and (b) we plot the E_1 field at $t' = 3746\ \omega_0^{-1}$ for simulations with either a modified high order solver or a hybrid Yee-FFT solver. Both solvers give nearly identical results and no evidence of NCI is observed in either case. In Fig. 5.8 (c) and (d) 2D plots of the electron density in the two cases are given. We also plot the line out of the on-axis E_1 fields for different time points in the boosted frame, as shown in Fig. 5.8 (e)–(h). As we can see from the comparisons, very good agreement between the results with these two solvers is obtained.

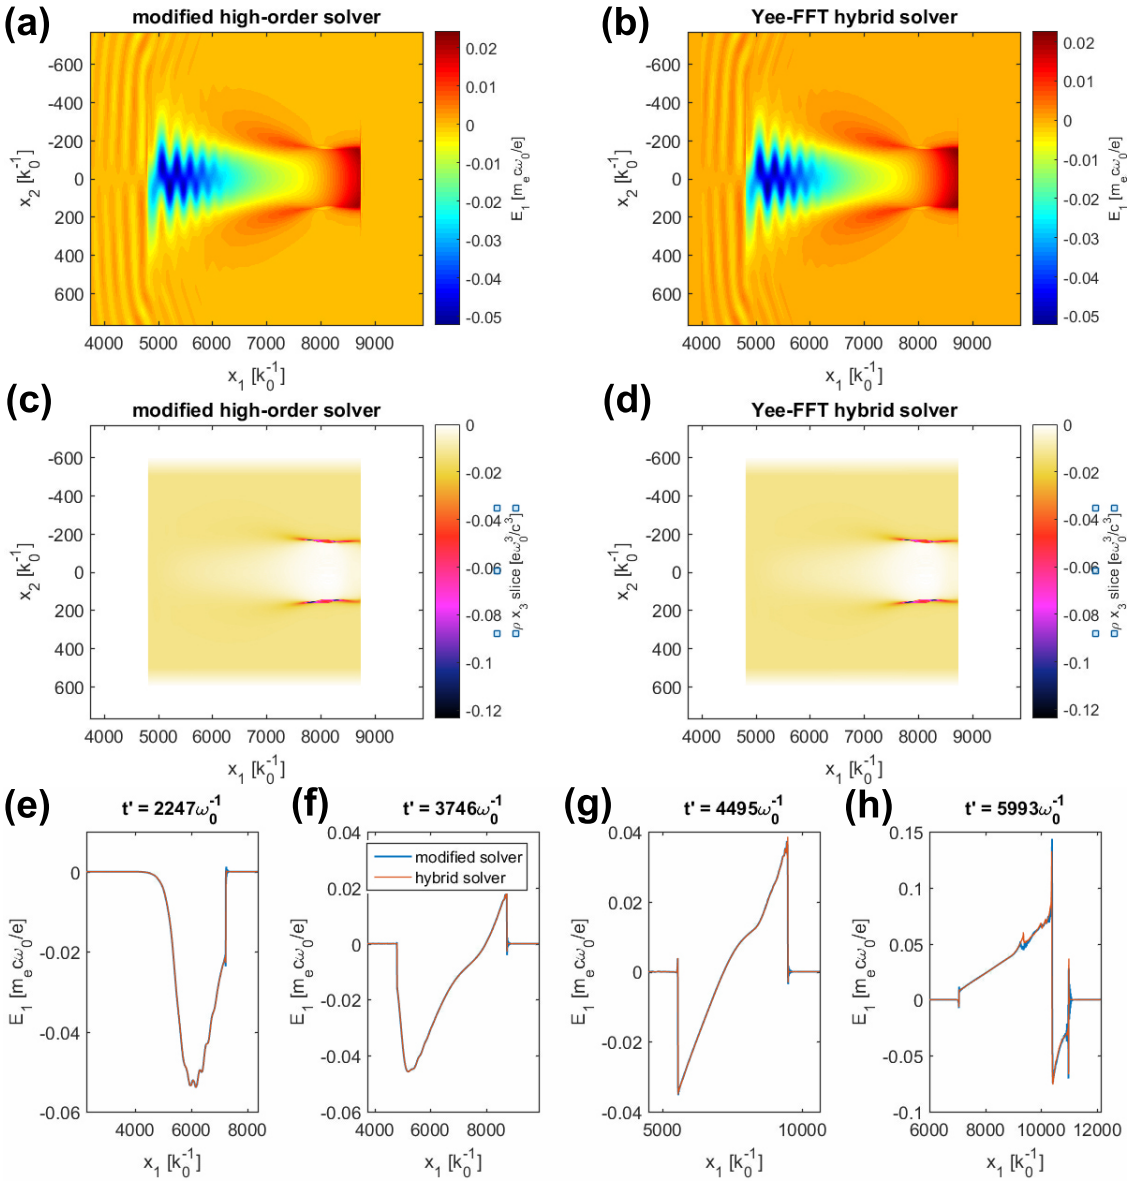

Figure 5.8: Comparison of simulations in the boosted frame between the customized high order solver and Yee-FFT hybrid solver. 2D plots of E_1 field at $t' = 3746\ \omega_0^{-1}$ for both solvers are shown in (a) and (b). The electron density profiles are shown in (c) and (d). (e) to (h) plot the on-axis lineouts of E_1 fields at different times.

Parameters	**Values**
Plasma	
density n_p	$8.62 \times 10^{-4} n_0 \gamma_b$
length L	$8.0 \times 10^4 k_0^{-1}/\gamma_b$
Laser	
normalized vector potential a_0	4.0
focal waist w_0	$153.0 k_0^{-1}$
pulse length τ	$86.9 k_0^{-1} \gamma_b (1+\beta_b)$
polarization	circular
Simulation setups	
cell size $\Delta x_{1,2,3}$	$0.1 k_0^{-1} \gamma_b (1+\beta_b)$
time step $\Delta t/\Delta x_1$	0.125
number of cells	$2048 \times 512 \times 512$
particle shape	quadratic
particle per cell	$(2, 2, 2)$
NCI elimination parameters	
Customized solver	
$[k_1]$ modification $(k_{1l}, k_{1u}, \Delta k_{\mathrm{mod,max}})$	$(0.1, 0.35, 0.01)$
lowpass filter (f_l, f_u)	(0.3,0.325)
Hybrid Yee-FFT solver	
$[k_1]$ modification $(k_{1l}, k_{1u}, \Delta k_{\mathrm{mod,max}})$	$(0.141, 0.24, 0.007)$
lowpass filter (f_l, f_u)	(0.3,0.35)

Table 5.5: Parameters for a 3D LWFA simulations in the Lorentz boosted frame using the customized high order solver and hybrid Yee-FFT solver. The laser frequency ω_0, wave number k_0 and the critical density $n_0 = m_e \omega_0^2/(4\pi e^2)$ in the lab frame are used to normalize the simulation parameters. The parameters of $[k_1]$ modification are normalized to $k_{g1} \equiv 2\pi/\Delta x_1$.

5.4 Summary

In this chapter, we have presented a new customized high-order FDTD solver combined with a current correction (such that Gauss's law remains satisfied) that effectively eliminates the NCI. The current is corrected and filtered by using a local FFT on each parallel partition when using domain decomposition. The customized higher order solver, and the corresponding current correction/filtering that is done locally on each partition, permits the systematic elimination of the Numerical Cerenkov Instability (NCI), while also permitting high parallel scalability in particle-in-cell codes without errors in the longitudinal fields. Using the theoretical framework we developed previously [45, 46] and illustrative PIC simulations, it is found that a high-order FDTD solver has similar NCI properties to that of a fully spectral solver or a hybrid Yee-FFT solver. By reducing the time step, the fastest growing $(\mu, \nu_1) = (0, \pm 1)$ NCI modes and $(\mu, \nu_1) = (0, 0)$ NCI modes can reside very close to the edge of the fundamental Brillouin zone. This enables the use of a lowpass filter on the current to effectively eliminate the NCI. For regular high-order FDTD solvers, highly localized NCI modes [which are part of the $(\mu, \nu_1) = (0, 0)$ modes] are seen in analogy to those observed in a spectral or hybrid Yee-FFT solver. These modes reside close to the physical modes in $\vec{k}$-space. Elimination of these modes can be achieved by a combination of applying reduced time step and creating a bump in the EM dispersion relation in k_1 space. This solver can be readily implemented in 2D/3D Cartesian and quasi-3D geometries contained within the existing framework of OSIRIS without the need to modify the boundary conditions in the transverse directions. We note that the boundary conditions in the $\hat{1}$ direction do not need to be changed since we can gradually reducing the order of the solver from 16th to 2nd order in the last 16 cells to match the boundary condition.

When the finite difference operators are modified, then the charge conserving current deposit must also be appropriately modified. We first deposit the current using the second order accurate charge conserving current deposit [23] in OSIRIS. The current is then Fourier transformed on each local partition, and then corrected, and filtered; it is then transformed back to real space for use in the field solver. The use of a current deposit that satisfies the

continuity equation for the higher order divergence operator is necessary such that Gauss' Law remains satisfied at each time step. We show that making such correction to the current will expand the range of cells over which the current for a particle is increased. Theoretically, a delta function for the current will extend to the entire simulation domain. However, the current falls below the double precision roundoff within a finite number of cells. Therefore, the current from a single particle is effectively localized. This permits using FFTs and the current correction and filtering for only the data on each parallel partition if the number of guard cells is properly chosen.

We have shown how the customized solver, together with its NCI elimination scheme, can systematically eliminate the NCI in a single drifting plasma. We have also shown how this scheme can be applied to relativistic shock simulation, with excellent NCI elimination achieved without sacrificing the parallel scalability of an FDTD EM-PIC code for problems with disproportionate number of cells in one direction. We have also shown the usefulness of the proposed high-order solver combined with local FFTs by conducting full 3D LWFA simulations in a Lorentz boosted frame.

CHAPTER 6

LWFA Lorentz boosted frame simulation

Now that the elimination strategies and techniques have been well discussed and developed in the previous chapters, in this chapter we discuss specifics on additional issues for performing LWFA simulation in the Lorentz boosted frame. In order for this chapter to be self-contained, we reuse some material from earlier chapters.

As mentioned in Chapter 1, when modeling LWFA physics using EM-PIC code, one has to resolve the smallest physical length of interest, in this case, the laser wavelength on the scale of a micro-meter, in order to ensure the evolution of the laser profiles is accurately simulated. On the other hand, the plasma column length that the laser driver is propagating through is on the centi-meter to meter scale. The disparity in these two critical physical lengths makes the simulation of LWFA very CPU time consuming, as the algorithm can only advance the laser propagation at the speed of micro-meter per time step. For instance, using a standard PIC code to study a 10 GeV stage in a nonlinear regime takes approximately 10–100 million core hours on today's computers. While computing resources now exist to do a few of such simulations, it is not possible to do parameter scans and convergence tests in full three-dimensions. Therefore, reduced models such as combining the ponderomotive guiding center with full PIC [24] for the wake or with quasi-static PIC [25, 26] are used for parameter scans. However, while these models are very useful, they cannot model full pump depletion distances and the quasi-static approach cannot model self-injection.

Recently, it was shown that by performing the simulation in an optimal Lorentz boosted frame with velocity v_b, the time and space scales to be resolved in a numerical simulation may be minimized [27, 29]. The basic idea is that in the boosted frame the plasma length (the laser propagation distance) is Lorentz contracted while the plasma wake wavelength

and laser pulse length are Lorentz expanded. The number of laser cycles is an invariant (assuming there is no or negligible reflected wave) so the necessary number of cells needed to resolve the laser is also an invariant while the cell size and hence time step are Lorentz expanded. The increase in time step and decrease in the plasma length lead to savings which scales with the factors $\gamma_b^2 = (1 - v_b^2/c^2)^{-1}$, as compared to a lab frame simulation using the so-called moving window [28]. Using such simulations, it has been shown that using a 1–3 PW laser one could generate a 10 GeV electron beam in a self-guided stage and a 50 GeV beam in a channel guided stage [32]. For these cases the savings can be larger than factors of 10^4. The boosted frame technique opens the possibility to study how the output electron energy of the LWFA scales to much higher energy when the laser energy is increased and the plasma density is decreased.

However, in the boosted frame LWFA simulations the NCI can be an issue. As discussed in previous chapters, the noise results from a numerical Cerenkov instability induced by the plasma drifting with relativistic speeds through the grid. According to the dispersion relation this numerical instability is attributed to the coupling between the wave particle (or beam) resonances with EM modes (including aliased modes) in the numerical system. The pattern of the instability in Fourier space can be found at the intersections of the EM dispersion relation of the solver used in the simulation algorithm, and the wave particle resonances (or more accurately, from coupling between EM and Langmuir modes).

In this chapter we mitigate this instability by using a multi-dimensional spectral Maxwell solver that greatly reduces growth rate of the numerical instability at the main NCI modes. In this case, the instability occurs only at high $|\vec{k}|$ modes which are far away from the physics of interest. As the EM dispersion curves for most finite-difference time domain (FDTD) solvers inevitably bends down (i.e., supports waves with phase velocities less than the speed of light) at high $|\vec{k}|$, a broad spectrum of additional numerical instability modes at the main beam resonance are found in these solvers. However, when using a spectral solver that spatially advances the EM fields in Fourier space, the NCI at the main beam resonance is greatly reduced. In addition, the pattern at the first space aliasing beam mode is found to indeed be located at high $|\vec{k}|$ values that are far away from the interesting physics. For

the spectral solver the numerical Cerenkov instability is located at a predicted pattern in $\vec{k}$ space so it can be conveniently eliminated by applying simple filters directly in $\vec{k}$ space.

We describe the development of a fully parallelized three-dimensional electromagnetic spectral PIC code called UPIC-EMMA that was rapidly built using components of the UCLA PIC Framework (UPIC) [58]. Here we use the word "spectral" to mean the fields are expanded using global basis functions. We used a truncated Fourier series and used FFTs to solve the field equations. In other respects, the code uses similar methods as the FDTD PIC codes, including interpolation of particle information to and from a grid [22]. We demonstrate in Chapter 3 that through the use of appropriate filters, Lorentz boosted frame simulations of LWFA at the optimum frame velocities can be carried out without limitations from the numerical Cerenkov instability. We show that a simple low pass filter with a hard cutoff at $|\vec{k}|$ works very well. This completely eliminates modes with $|\vec{k}|$ above a selected value. Meanwhile, it is not as easy to use such a filter in $|\vec{k}|$ space using a FDTD solver (and such solvers have instabilities at lower $|\vec{k}|$). In Chapter 3 we likewise showed how to eliminate the main NCI modes.

We have benchmarked UPIC-EMMA by comparing the 2D and 3D simulation results of LWFA in Lorentz boosted frames with the corresponding OSIRIS lab frame simulations. Good agreement is found between the OSIRIS lab frame simulations, and UPIC-EMMA boosted frame simulations, in both linear, and nonlinear regimes. We also compare UPIC-EMMA simulations for different values of γ_b and excellent agreements are found.

It is worth noting that the main purpose of this chapter is to describe UPIC-EMMA, as well as to illustrate the detailed simulation setups for the LWFA Lorentz boosted frame simulations using UPIC-EMMA. Much of the material is taken from [71]. However, at the time [71] was published, we were not fully aware of the effect of the main NCI modes, which when using the spectral solver has a highly localized pattern with a growth rate one order of magnitude smaller than the first spatial aliasing modes. In the simulations presented in the remainder of this chapter, we have not eliminated the main NCI modes. The subtle differences between simulations with and without the elimination of the main NCI modes

will be discussed in section 6.4, and was discussed in Chapter 3.

The remainder of this chapter is organized as follows. In section 6.1 we review the firsts spatial NCI modes induced by relativistic drift. In section 6.2, we describe the development of UPIC-EMMA, and how using the algorithms in UPIC-EMMA can eliminate the instability induced by relativistic plasma drift. In section 6.3, we discuss details of LWFA Lorentz boosted frame simulations using UPIC-EMMA. In section 6.4, we benchmark UPIC-EMMA results with different γ_b, and with OSIRIS lab frame simulations. A summary is given in section 6.5.

6.1 NCI in LWFA boosted frame simulation

The numerical Cerenkov instability induced by relativistic plasma drift has been extensively discussed in previous chapters. In a PIC system, when the plasma is drifting relativistically, the velocity of the drifting particles can be equal to the component of (be in resonance with) the phase velocity of the main EM mode along the drift direction. In addition, its aliased modes can always be in resonance with the EM modes. Due to the nature of wave-particle resonances, the numerical instability occurs at the intersections of the beam resonances and EM modes determined by the Maxwell solver used in the simulation. By carefully choosing the Maxwell solver, the instability pattern can be manipulated so that mitigation can be achieved. As discussed in Chapter 2 and 3, when a spectral solver is used, the growth rate of the main beam resonance is greatly reduced and becomes highly localized. As a result, the instability is dominated by the aliased resonances, and the fastest growing modes are the first spatial aliases. These resonances reside at high $|\vec{k}|$ in Fourier space far away from the important physics.

Since the instability pattern is found near the intersections of the EM modes and beam resonance, we can obtain a simple analytical expression for the instability pattern in the limit $\Delta t \to 0$ (which leads to $[\omega] = \omega$). Under this assumption, in the 2D scenario the

equation for the EM dispersion curves in the spectral solver is

$$\omega^2 \approx k_1^2 + k_2^2$$

And the equations for the beam resonances are

$$\omega + \mu\omega_g = \beta(k_1 + \nu_1 k_{g1})$$

where $\beta \equiv v/c$. Defining $\xi \equiv \beta\nu_1 k_{g1} - \mu\omega_g$, we can obtain the expressions for the intersections as

$$(1-\beta^2)k_1^2 + k_2^2 - 2\beta\xi k_1 - \xi^2 = 0 \tag{6.1}$$

Note that the fastest growing modes for the instability pattern are at the $\mu = 0$ and $\nu_1 = \pm 1$ resonances. In the limit of interest $\beta \to 1$, we obtain

$$k_2^2 \mp 2k_1 k_{g1} - k_{g1}^2 = 0 \tag{6.2}$$

In figure 6.1 (a) we plot Eq. (6.2) for $\mu = 0$ and $\nu_1 = \pm 1$. Note the "ring" pattern of the instability, which crosses the $k_2 = 0$ axis near the point $(\pm k_{g1}/2, 0)$; therefore this mode is located at high $|\vec{k}|$ values which are far away from the region of interesting physics. Therefore, the numerical Cerenkov instability can be effectively eliminated in some cases if the fastest growing modes ($\mu = 0, \nu_1 = \pm 1$) are suppressed in the Maxwell solver. Note in Chapter 2 and 3, we discussed how to eliminate additional modes.

In the simulations, we identify the unstable modes in Fourier space using the approximate expression Eq. (6.2). We then apply filters with specific masks which multiply the undesired modes by zero. In figure 6.1 (b) we plot the "ring-shaped" filter used in some of the two-dimensional simulations for testing the instability mitigation. We put "ring" in quotes because it is not a true ring but rather a range between two parabolas. We also used a low pass filter with a hard cut-off. We filled the simulation box with neutral plasma drifting relativistically at $\gamma = 14000$ in the x_1 direction, and ran cases without a filter, with the "ring" filter, and the low pass filter with a hard cutoff. As seen from figure 6.1 (c), these filters efficiently suppresses the instability modes at $\mu = 0, \nu_1 = \pm 1$ in E_2. Therefore, the

mitigation of the instability using this filter shows the flexibility and efficiency of a spectral solver in being able to pinpoint the suppression of the unphysical modes in PIC simulations while leaving the modes near the interesting physics completely unaffected.

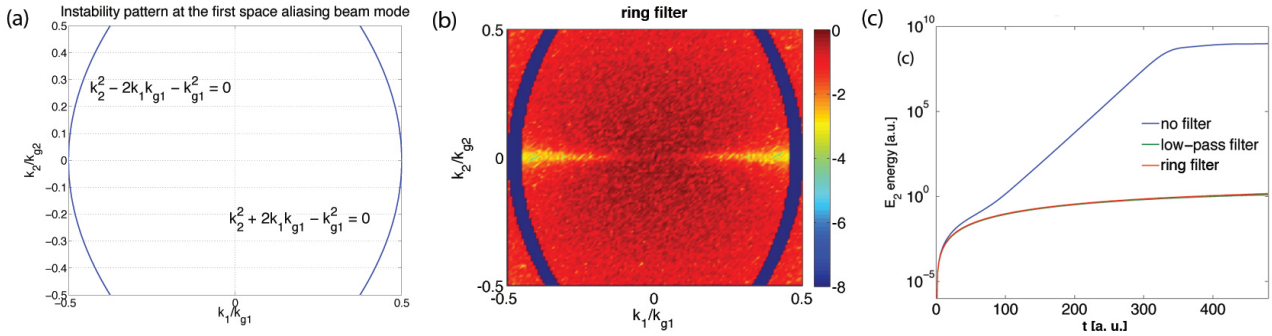

Figure 6.1: (a) shows the analytical expression for the $\mu = 0$, $\nu_1 = \pm 1$ mode of numerical Cerenkov instability for the 2D spectral solver in (k_1, k_2) plot; (b) shows the "ring-shaped" filter applied in the 2D spectral solver; and (c) shows the E_2 energy evolutions for various simulation setups.

6.2 EM-PIC code with spectral solver

As mentioned in the section 6.1, an EM-PIC code with a spectral solver has superior properties in suppressing the numerical Cerenkov instability induced by a relativistic plasma drift. They also have superior properties with respect to numerical dispersion errors and noise. We will briefly explain the algorithm of a spectral EM-PIC code, as well as discuss the challenges in optimizing the performances of such a code in the following paragraph.

Spectral PIC codes have a long history [20, 22]. However, despite their advantages in better accuracy and less noise, they are not currently as widely used, partly because implementing boundary conditions and sustaining high parallel scalability are not as straight forward in spectral PIC code. A spectral EM-PIC code has the same basic flow chart as an FDTD PIC code. In a spectral EM-PIC code both the charge and current are deposited on the mesh from the particles; the forces exerted on the particles are interpolated from

the mesh points, and particles are advanced using the Lorentz forces. The main difference between the spectral PIC code and FDTD PIC code is the solver used to advance the electromagnetic field, and that in the spectral PIC code all field quantities, including the charge and current densities, are defined at the same locations on a cell (no Yee mesh [21] is needed). In a spectral PIC code the charge and current are directly deposited, and a strict charge conserving current deposit is not needed because Gauss's law is solved at each time step using the charge density [22]. This gives the longitudinal part of the electric field. The longitudinal component of the magnetic field is set to zero at each time step. Faraday's law and Ampere's law are used to advance the transverse electric and magnetic fields forward in time. Note that because Gauss's law is solved for directly at each time step, a charge conserving current deposit or Boris correction to the longitudinal component of the electric field is not required to maintain that Gauss's law is satisfied. The equation for the longitudinal component of electric field $\vec{E}_L$ becomes:

$$\vec{E}_L(\vec{k}) = -4\pi\rho(\vec{k})\frac{i\vec{k}}{k^2} \tag{6.3}$$

and the transverse electric field and magnetic field are leap-frogged forward in time using Faraday's and Ampere's law:

$$\frac{\partial\vec{E}_T(\vec{k})}{\partial t} = ic\vec{k}\times\vec{B}(\vec{k}) - 4\pi\vec{j}_T(\vec{k}) \qquad \frac{\partial\vec{B}(\vec{k})}{\partial t} = -ic\vec{k}\times\vec{E}_T(\vec{k}) \tag{6.4}$$

where the transverse component of the current is:

$$\vec{j}_T = \vec{j} - \frac{\vec{k}(\vec{k}\cdot\vec{j})}{k^2} \tag{6.5}$$

We also multiply $\rho(\vec{k})$ and $\vec{J}(\vec{k})$ by a shape function $S(\vec{k}) = \exp(-|k|^2a^2/2)$ where a is the particle size. The fields are also multiplied by this shape function then interpolated to the particles [20].

In addition, just as in a FDTD code, the particle positions and velocities (and correspondingly the charge and current densities) are defined at half integer values in time with respect to each other. If positions are defined at whole time steps and velocities (momentum) at half integer values, then the longitudinal and transverse components of the electric field

are defined at whole time steps (when particle positions are defined) and the magnetic field is defined at half-integer values. The current at the half integer time are time centered by averaging the new and old positions during the deposit. Once the fields are transformed back from $\vec{k}$ space to real space then the particles can be pushed. The particle push is identical to that of a FDTD except for the interpolation of the forces because all field quantities are defined at the same locations in a cell.

There are no dispersion errors due to the spatial grid. However, in our implementation where we use second order leap frog scheme for the time advance there are errors due to the finite time step. While a FDTD code (that uses two cells to define the derivative) describes the $[k]_i$ operator to $\mathcal{O}(\Delta x_i)^2$, the spectral code has a $[k]_i$ operator with the accuracy exceeding $\mathcal{O}(\Delta x_i)^N$ where N is the number of grid points. However, this improved accuracy arises because the equivalent finite difference operator is non-local [37]. This may be an issue when plasma or laser with very sharp gradient exists. We note that in an FFT based code $[\vec{k}] = \vec{k}$ can be replaced with $[\vec{k}]$ of any form. We use this feature in a test-bed code to study the effects of local v.s. global operators. Both a spectral and a FDTD code effectively truncate the highest $|k_i|$ to $\pi/\Delta x_i$. In addition, when including time step errors, the numerical dispersion of a spectral PIC code is superluminal, while that of the FDTD code is subluminal. As we discuss elsewhere in this chapter, the more accurate and superluminal aspect of the EM dispersion relation provided by the spectral solver (together with the simple filters) is crucial for eliminating the fastest growing modes of numerical Cerenkov instability. Others have discussed using the PSATD method to get perfect dispersion in vacuum [41]. However, we do not find this to be advantageous regarding NCI elimination. The corresponding Courant condition in 2D and 3D are (for the square and cubic cells) [20]:

$$\Delta t_{2D} = \frac{2}{\sqrt{2}\pi c} \qquad \Delta t_{3D} = \frac{2}{\sqrt{3}\pi c} \tag{6.6}$$

A spectral PIC code is also distinguished from a FDTD code in the way it is parallelized. For the field solver, the simulation box is usually partitioned in one dimension in 2D, and two dimensions in 3D, so that each processor holds global information in the dimension to be transformed. As a result, a parallel spectral PIC code requires a fast parallel transpose

routine for efficient execution of the FFT in multi-dimensions. The parallel spectral PIC code scales well on parallel computers if the problem size is large enough [22]. Scaling stops when the all-to-all transpose used in the FFT becomes latency dominated, which depends on the network being used. In many cases the decomposition for the particles is the same as that for the fields although this does not have to be the case. Recently, Vay et. al. have proposed solving for the transverse fields within each domain separately to reduce some parallel communication and showed some success in LWFA boosted frame simulation [54]. However, this technique may lead to errors in the longitudinal fields.

We have developed a multi-dimensional EM-PIC code called UPIC-EMMA that uses a spectral field solver. This code was rapidly put together using components provided by the UPIC Framework, a PIC framework with spectral solvers developed at UCLA [58]. UPIC-EMMA is fully relativistic and fully parallelized. Inherited from the UPIC Framework, UPIC-EMMA is coded in layers for convenient extension with different programming styles. The lowest layers are written in Fortran77 for high performance. They can be easily extended to many other languages. On top of this layer exists a library of Fortran90 wrapper functions which hide the complexity of the Fortran77 layer and that provide simpler arguments which enables strict type checking. The code separates the physics procedures from the communication, and utilizes the Message-Passing Interface (MPI) for parallel processing. In addition, a multi-tasking library was implemented to enable mixed multi-tasking and MPI messaging, where multi-tasking is used on a shared memory node with multiple cores, and message-passing is used between such nodes [58]. UPIC-EMMA also features 3D load balancing where the fields and particles use different partitions.

6.3 LWFA Simulations in the Lorentz boosted frame

In section 6.1 we described general issues regarding the numerical instability that arises when a plasma drifts near the speed of light. In this section we describe some details regarding issues specific to modeling LWFA in a Lorentz boosted frame. We describe issues related to numerical dispersion in the lab frame, in the boosted frame, and in transforming from the

boosted frame back to the lab frame for comparison. We also discuss the moving antenna and interactions between the laser and the drifting plasma boundary.

6.3.1 Numerical dispersion errors for the laser

As mentioned in the introduction, one of the first obstacles in modeling LWFA in a boosted frame is to mitigate the numerical Cerenkov instability. For the FDTD PIC code which uses a combination of a Yee solver together with the momentum conserving field interpolation scheme, it is useful to choose the optimal time step $\Delta t \approx \Delta x_1/2$, where the $\hat{1}$-direction is the plasma drifting direction, to minimize the numerical Cerenkov instability growth rate [40, 45, 53]. The need to use this time step eliminates the flexibility in tuning the time step to minimize numerical dispersion errors for the laser. In figure 6.2 we present the error in the group velocity of an EM wave on a grid in the 2D scenario (we let $\Delta x_1 = \Delta x_2$). Note that for the Yee, and Karkkainen solvers [59] which were discussed in Ref. [40, 45], the most accurate dispersion relation occurs at their Courant Limit, but not the corresponding optimal time step at $\Delta t \approx \Delta x_1/2$ (for momentum conserving field interpolation). On the other hand, for a spectral PIC code the instability mitigation does not rely on the relation of grid sizes and time step. In particular, the EM dispersion relation can be made arbitrarily accurate by reducing the time step [see figure 6.3 (a)]. As we showed in Chapter 3, this also greatly reduces the growth rate of the main NCI modes, and moves them away from the physical modes near $\vec{k} = \vec{0}$. In general when simulating relativistically drifting plasma, a spectral PIC code can provide more accuracy and flexibility over the FDTD PIC code with respect to numerical dispersion in the simulated frame.

6.3.2 Lorentz transform of boosted frame data

While numerical dispersion errors exist when using a finite size grid in vacuum, here we show that when modeling the LWFA in the Lorentz boosted frame, these errors in the boosted frame are not necessarily an issue when the results are transformed back to the lab frame. While the value for γ_b in the boosted frame can be arbitrary, the speed up increases as

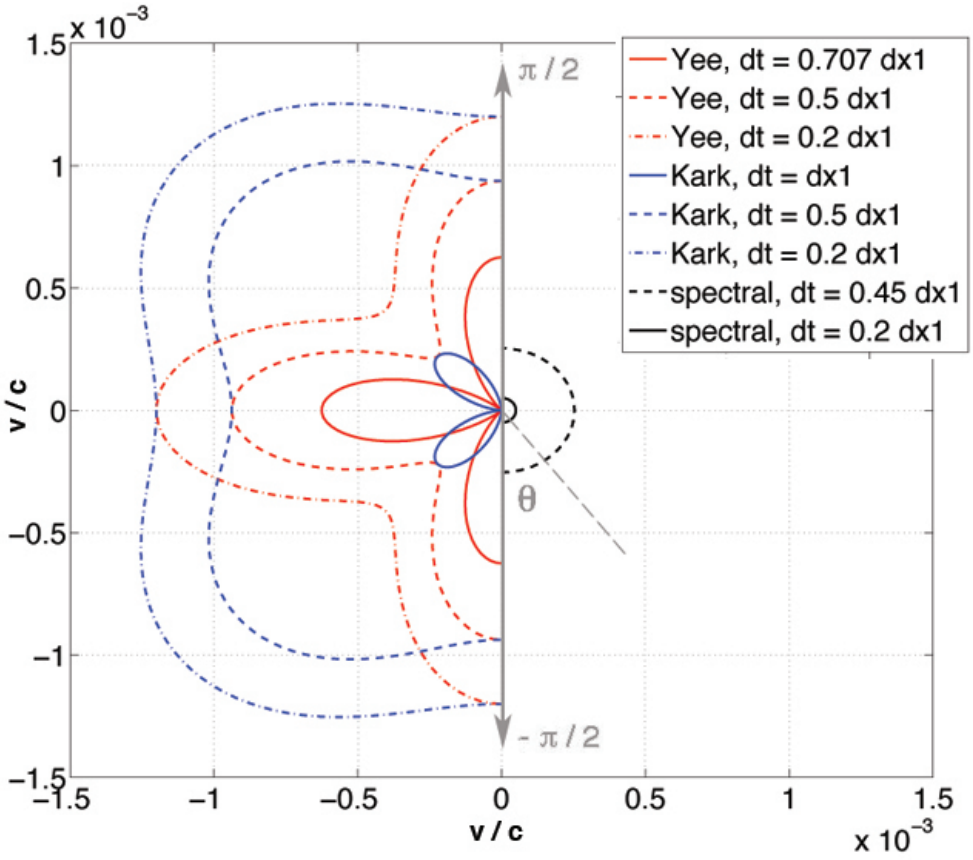

Figure 6.2: The plot shows the errors in the group velocity defined as $v_g - 1$ of the 2D EM dispersion relation for various cases. Defining $\theta = 0$ to be the laser propagating direction, this plot shows the propagation angle in $(-\pi/2, \pi/2)$. If the error $(v_g - 1)$ is larger than zero, its corresponding point is in the right side of the vertical axis, and vice versa. The group velocity is calculated for the $k_0 = 1.0$ mode while we are using $k_0\Delta x_1 = k_0\Delta x_2 = 0.1$ for the calculation.

γ_b is increased. However, choosing $\gamma_b \approx \gamma_w$, where γ_w is the phase velocity of the wake, is generally optimum because in this frame the plasma length and the laser pulse length are nearly matched. When the laser and plasma frequency are comparable, each is resolved similarly, i.e., there is no over-resolution of either the laser wavelength or plasma wavelength. In the boosted frame, the length of the plasma contracts by γ_b, the electron and ion mass are both γ_b times heavier, the plasma density is γ_b times larger, and the corresponding plasma frequency is a Lorentz invariant. As for the laser, there is a $\gamma_b(1 + \beta_b)$ stretch to the pulse length, while the Rayleigh length contracts by γ_b. Therefore, while the pulse waist does not change, the effective spot size at the rear of the pulse increases by a factor of $\gamma_b^2(1+\beta_b)$. Hence for sufficiently large γ_b an antenna is needed to launch the laser from the laser pulse waist that is moving backwards [60, 61]. The antenna is usually placed at the plasma boundary (see section 6.3.3 for details).

In the lab frame simulation, a moving window which only models the region of interest around the laser is often used to reduce the simulation box size. Implementation of a moving window is challenging in a spectral PIC code due to the non-local nature of the field solver which necessitates knowledge of boundary condition at both of the moving boundaries. However, the relative range of x_1 and t contracts when Lorentz transforming the data of interest from lab frame to boosted frame [33, 62]. If γ_b is appropriately chosen, in this frame the length of the plasma column is of the same order as the laser pulse length [16]. As a result, for $\gamma_b \sim \gamma_w$ it is feasible to conduct the boosted frame simulation without the moving window. We discuss a moving window in the boosted frame that follows the drifting plasma in the next chapter.

In LWFA lab frame simulations, an EM wave with frequency ω_0 is incident on a stationary plasma slab. This leads to reflected and transmitted waves, each having the incident frequency. Their wave numbers are determined from the dispersion relation in vacuum (reflected wave), and in plasma (transmitted wave). In a simulation the same physics occurs except the EM wave now satisfies the numerical dispersion relation in vacuum and plasma. In the boosted frame there are still reflected and transmitted waves, except in this case the incident wave, reflected wave, and transmitted wave each have different frequencies. Fur-

thermore, numerical issues can lead to some subtle effects. An effective method to identify the frequencies of the reflected and transmitted waves is to use an (ω, k) diagram, which was previously used in studying the radiation generated from ionization fronts [70]. At the plasma boundary $z_0 = -v_b t$, the phase of each wave $\phi = kz - \omega t = -(kv_b + \omega)t$ must be the same, otherwise the continuity of fields cannot be satisfied at every instant in time. This leads to

$$k_i v_b + w_i = -k_r v_b + \omega_r = k_t v_b + \omega_t \tag{6.7}$$

where i, r, t correspond to incident, reflected, and transmitted waves respectively. For example, if $v_b = 0$ then $\omega_i = \omega_r = \omega_t$. If the incident and reflected waves obey the vacuum dispersion relation $\omega = k$ then

$$\omega_r = \frac{1 + \beta_b \omega_i}{1 - \beta_b} \tag{6.8}$$

which can also be obtained from a double Lorentz transformation. In a Lorentz boosted frame the plasma is drifting but $\omega_i = \omega_0$ is Lorentz transformed to ω_i' and we want ω_r' and ω_t' [where the $(')$ sign refers to the boosted frame variables]. In this frame

$$\omega' + k' v_b = \omega_0' + k_0' v_b \tag{6.9}$$

where ω' can be either the reflected or transmitted wave. The constant $\omega_0' + k_0' v_b$ is obtained by Lorentz transforming ω_0 and k_0 into the boosted frame: $\omega_0' = \gamma_0(\omega_0 - k_0 v_b)$ and $k_0' = \gamma_0(k_0 - v_b \omega_0/c^2)$. Therefore, $\omega' + k' v_b = \omega_0/\gamma_b$ regardless of the relationship between ω_0 and k_0. In a real system $\omega_0 = k_0$ although numerical errors in the dispersion relation do not alter the constant ω_0/γ_b. Therefore, the reflected and transmitted waves must fall along the line $\omega' = -k' v_b + \omega_0/\gamma_b$ in (ω', k') space (here we are ignoring the aliasing modes). In addition, they must also fall on the dispersion curves for light in a plasma [45]

$$[\omega]^2 = [k]^2 c^2 + \frac{\omega_p'^2}{\gamma_b} S^2 \frac{[\omega] - [k] v_b}{\omega - k v_b} \tag{6.10}$$

or in vacuum

$$[\omega]^2 = [k]^2 c^2 \tag{6.11}$$

on the grid where we assume $S = S_{j3} = S_{E3} = S_{B2}$, and $\omega_p'^2/\gamma_b = 4\pi e^2 n_0'/m_e$ is Lorentz invariant where n_0 is the lab frame density, e is the electron charge, and m_e is the electron rest mass. The reflected and transmitted waves lie at the intersection between Eq. (6.9), and Eq. (6.10) in plasma or Eq. (6.11) in vacuum. This is shown in figure 6.3 (a) for a case where $\Delta t \approx 0.5\Delta x_1$, $\omega_0/\omega_p \approx 30$, and $\gamma_b = 8.0$. The line $\omega' = -v_b k' + \omega_0/\gamma_b$ and the dispersion curve for a real plasma (black dashed lines), a FDTD Yee solver (red lines), and a spectral solver (green lines for inside the plasma; magenta lines for in vacuum) are shown. The vacuum dispersion relation is plotted for the spectral solver in the upper left quadrant for the reasons given in the next paragraph. In figure 6.3 (b), we have expanded the region in (ω, k) space near the origin to illustrate the frequency and direction of the transmitted wave which does not depend strongly on the solvers used. When $\gamma_b = \gamma_w = \omega_0/\omega_p$ then $\omega_t' = \omega_p = \omega_p'/\sqrt{\gamma_b}$ and $k_t' = 0$. If $\gamma_b > \gamma_w$ then ω_0' would be negative and the phase velocity and group velocity of the transmitted wave would be negative; however, since $|v_{gt}'| < |v_b|$ the transmitted wave would still be in the plasma.

Figure 6.3 (a) also illustrates that numerical errors in the dispersion relation effect the location of the reflected wave. In a real system where $\omega' = ck'$ in vacuum and $\omega'^2 = \omega_p'^2/\gamma_b^2 + c^2 k'^2$ in the plasma, then the reflected wave would occur where $\omega' = -v_b k' + \omega_0/\gamma_b$ intersects the vacuum curve, i.e., at $\omega' = \omega_b\gamma_b(1+\beta_b) \sim 2\omega\gamma_b$, which is larger than the largest ω' in the fundamental Brillouin zone. However, for the numerical dispersion curves shown in figure 6.3 (a), the reflected wave resides at the intersection with the plasma dispersion relation in the lower right quadrant for the FDTD solver or with the vacuum dispersion relation in the upper left quadrant for the spectral solver. For the FDTD case, the reflected wave has a negative phase and group velocity. However, the group velocity is less than v_b so the reflected wave propagates backwards while staying inside the plasma. For the spectral solver the group velocity is slightly larger than the speed of light so it resides outside the plasma. The predicted locations of the transmitted and reflected waves are confirmed in an OSIRIS (FDTD) simulation. This is seen in figure 6.3 (c), in which the ω' and k' spectrum is plotted from a simulation for parameters identical to those used to generate the theoretical plot in figure 6.3 (a). Strong signals are seen at the predicted locations. For cases of

interest the reflection coefficient is small [the reflected signal is significantly smaller than the transmitted signal in figure 6.3 (c)] so the unphysical mode is not energetically important, and it does not complicate the physics.

We have also investigated the invariance of transforming results back to the lab frame based on the numerical dispersion relations. It has been demonstrated that good agreements can be found when the boosted frame data are transformed back to the lab frame [60]. However, we note that solving Maxwell's equations on a grid using discrete time steps is not strictly Lorentz invariant. For example, the group velocity of light in vacuum for a spectral solver is greater than the speed of light, and it depends on $\omega\Delta t$. Nevertheless, when carrying out LWFA (or other) simulations in a boosted frame, the results are transformed back to the lab frame using the Lorentz transformations. This is assumed to be reasonable if one is looking at modes which are properly resolved.

As noted earlier, when γ_b is chosen near γ_w there is a balance between the laser pulse length and the plasma length. In addition, for $\gamma_b \approx \gamma_w$ the transmitted wave $k' \sim 0$, and errors in the boosted frame due to the finite cell size are minimized. In figure 6.3 (d) we show that when v'_ϕ and v'_g for the transmitted wave are Lorentz transformed back to the lab frame using the velocity addition formulas,

$$\beta_\phi = \frac{\beta'_\phi + \beta_b}{1 + \beta'_\phi \beta_b} \qquad \beta_g = \frac{\beta'_g + \beta_b}{1 + \beta'_g \beta_b} \tag{6.12}$$

where β_ϕ and β_g are the phase and group velocity normalized to c, that the numerical errors are nearly absent for sufficiently large γ_b. The values for v'_ϕ and v'_g are calculated from the linear dispersion relation, where $\omega\Delta t = 0.5 k\Delta x_1$ is given and $k\Delta x_1 = 0.2$ is from the dispersion relation in the lab frame. In the boosted frame $\Delta t' = \gamma_b(1 + \beta_b)\Delta t$ and $\Delta x'_1 = \gamma_b(1 + \beta_b)\Delta x_1$. According to the plot, for $\gamma_b = 1$ there are clear numerical errors; however, for $\gamma_b \geq 5$, the numerical errors are minimized.

These results can be understood as follows. In a Lorentz boosted frame where $\beta_b = \beta_w \equiv (1 - \gamma_w^{-2})^{1/2}$, the group velocity $\beta'_g \to 0$, while the phase velocity $\beta'_\phi \to \infty$ in the numerical system, which when substituted back to Eq. (6.12) leads to

$$\beta_\phi = 1/\beta_w \qquad \beta_g = \beta_w \tag{6.13}$$

which are the accurate values for a continuous system. In addition, writing $\beta = \overline{\beta} + \delta_\beta$, and $\beta' = \overline{\beta}' + \delta'_\beta$, where $\overline{\beta}$ and $\overline{\beta}'$ corresponds to the correct values in the lab and boosted frame, and defining

$$\overline{\beta}_{\phi,g} = \frac{\overline{\beta}'_{\phi,g} + \beta_b}{1 + \overline{\beta}'_{\phi,g}\beta_b} \tag{6.14}$$

we can obtain the expressions of the error $\delta_\beta \equiv \beta_{\phi,g} - \overline{\beta}_{\phi,g}$ as

$$\delta_\beta = \frac{1}{(1 + \overline{\beta}'\beta_b/\delta'_\beta + \beta_b)(1 + \overline{\beta}'\beta_b/\delta'_\beta)\gamma_b^2} \tag{6.15}$$

Note the γ_b^2 in the denominator indicates that when γ_b is sufficiently large, the errors in velocity when transformed back to the lab frame will be small for all waves. We also note that the arguments going from Eq. (6.12) to Eq. (6.15) hold for any velocity including those of the particles. This indicates that if we choose the γ_b large enough that we would obtain more accurate results compared to a simulation done in the lab frame (with typical cell sizes and time steps).

6.3.3 Moving antenna

As discussed in [33, 60], the effective spot size of the laser increases by a factor of $\gamma_b^2(1+\beta_b)$ because the Rayleigh length of the laser contracts by γ_b and the pulse length expands by $\gamma_b(1+\beta_b)$. To prevent the need for using a simulation box with transverse size $\sim \gamma_b^2$ times that in needed in the lab frame, we utilize a thin slice of grids at the plasma boundary (where the laser beam waist resides) as an antenna to drive the laser pulse into the plasma [60, 61]. The antenna is moving together with the plasma boundary [see figure 6.4].

The EM field in the moving antenna as a function of $\vec{x}$ and time t can be derived as follows. For instance, for a laser linearly polarized in the $\hat{2}$ direction, the expression for the electric field $E_2(\vec{x}, t)$ of a Gaussian pulse in the lab frame can be expressed as:

$$E_2(x_1, x_2, x_3, t) = \frac{E_0 W_0}{W(x_1)} \exp\left[-\frac{x_2^2 + x_3^2}{W^2(x_1)}\right] \exp\left[-\frac{2(x_1 - ct)^2}{\sigma_s^2}\right]$$
$$\exp\left[ikx_1 + ik\frac{x_2^2 + x_3^2}{2R(x_1)} - i\arctan\frac{x_1}{x_R}\right] \exp(-i\omega t)$$

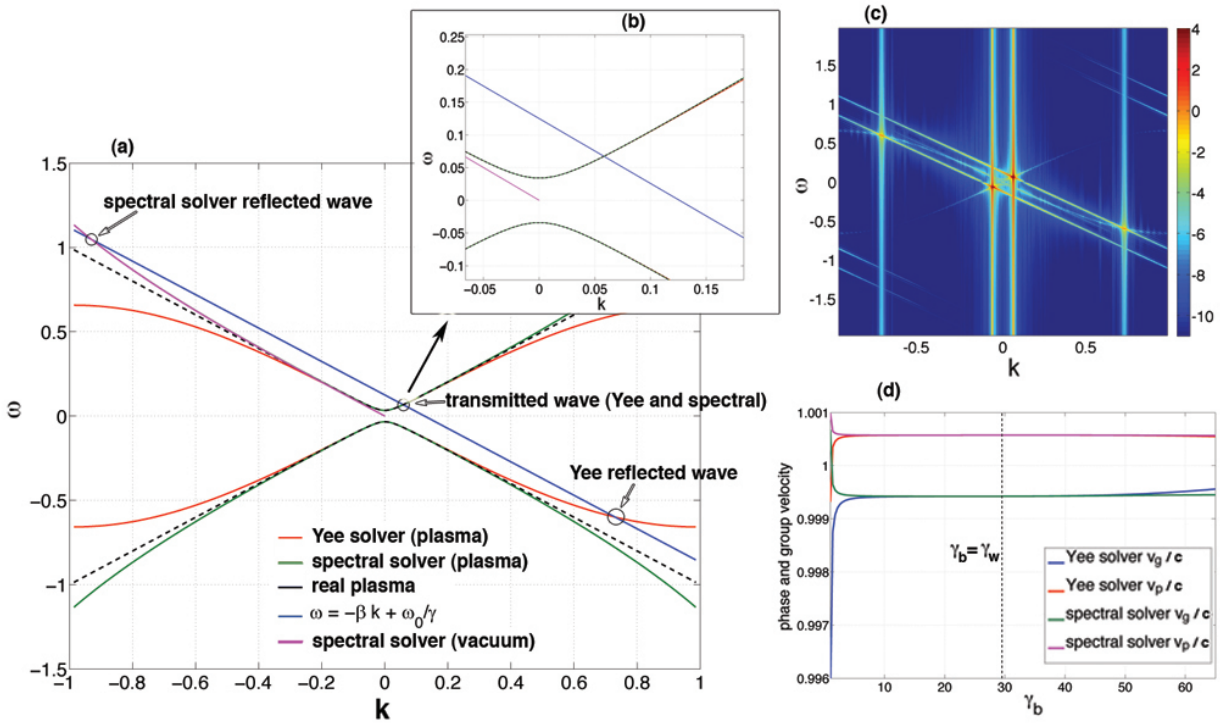

Figure 6.3: (a) shows the intersections of the line $\omega = -v_b k + \frac{\omega_0}{\gamma_b}$ and various EM dispersion curves, while in (b) we magnified the region near the origin; (c) shows an example of the E_3 spectrum of a 1D LWFA boosted frame simulation with the Yee solver. The hot spots in (c) show where the transmitted and reflected waves are, and agrees with the prediction in (a). (d) shows the dependence of the transformed phase and group velocity of the EM waves in the plasma with γ_b. The phase and group velocity converges quickly as γ_b increases from 1.

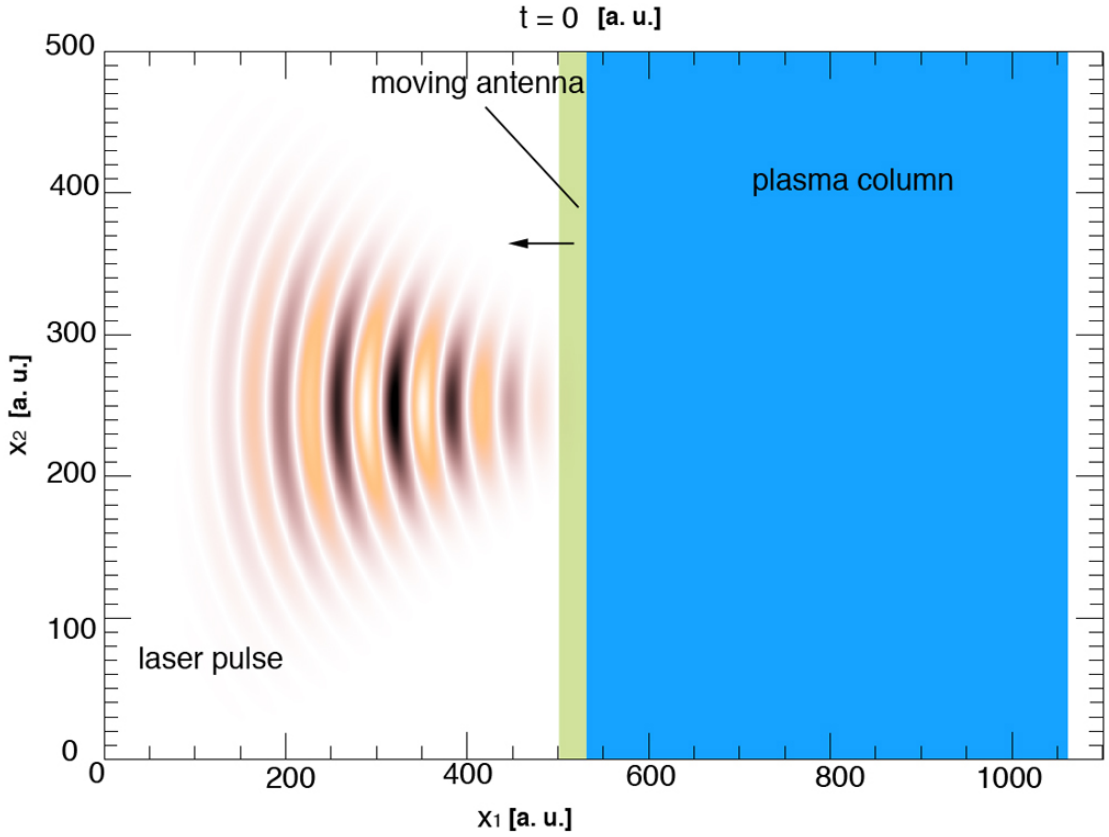

Figure 6.4: UPIC-EMMA simulation setup for LWFA boosted frame simulation. The blue block is the plasma column; the green slice is the moving antenna at $t = 0$. The laser is launched via the moving antenna (moving together with the plasma column boundary at $v = -\beta_b$) by initializing the appropriate curent in the green slice which has a typical width of $\lambda/2$. The laser is likewise plotted for $t = 0$. Note when the laser is launched via the antenna, only the area within the antenna is initialized.

with

$$W(x_1) = W_0\sqrt{1+\frac{x_1^2}{x_R^2}} \qquad R(x_1) = x_1\left(1+\frac{x_R^2}{x_1^2}\right) \qquad x_R = \frac{\pi W_0^2}{\lambda}$$

where E_0 is the amplitude, W_0 is laser pulse waist, σ_s is the laser pulse length, ω and k are the laser frequency and wavenumber, and x_R is the laser Rayleigh length. A similar expression holds for the magnetic field $B_3(\vec{x},t)$. After Lorentz transforming, we can readily obtain the new expression of the laser pulse in the boosted frame

$$\begin{aligned} E_2'(\eta, x_2', x_3', t') = \; & \frac{E_0' W_0}{W'(\eta)}\exp\left[-\frac{x_2'^2 + x_3'^2}{W'^2(\eta)}\right]\exp\left[-\frac{2[\eta-(1+\beta)ct']^2}{\sigma_s'^2}\right] \\ & \exp\left[ik'\eta + ik'\frac{x_2'^2+x_3'^2}{2R'(\eta)} - i\arctan\frac{\eta}{x_R'}\right]\exp(-i\omega' t') \end{aligned} \tag{6.16}$$

where

$$\eta = x_1' + \beta_b ct' \qquad \sigma_s' = \gamma_b(1+\beta_b)\sigma_s \qquad E_0' = \frac{E_0}{\gamma_b(1+\beta_b)} \tag{6.17}$$

$$k' = \frac{k}{\gamma_b(1+\beta_b)} \qquad \omega' = \frac{\omega}{\gamma_b} \qquad x_R' = \frac{x_R}{\gamma_b} \tag{6.18}$$

$$W'(\eta) = W_0\sqrt{1+\frac{\eta^2}{x_R'^2}} \qquad R'(\eta) = \frac{\eta}{1+\beta_b}\left(1+\frac{x_R'^2}{\eta^2}\right) \tag{6.19}$$

In the spectral code, the transverse and longitudinal components of the fields are solved for separately [22]. Therefore, on the antenna we set $\rho = 0$ so there are no longitudinal fields on it. When launching a laser from the antenna, we assign current (in the direction of the laser polarization direction) at every point inside the antenna such that $\vec{E}$ has the desired form and polarization. The other components and the magnetic field follow naturally from the Maxwell field solver. The antenna has a finite width of around $\lambda/2$ where λ is the wavelength of the laser in vacuum to eliminate any backward propagating signal. The current for generating the laser is added after the current is deposited for all the particles in the system.

The moving antenna implemented in UPIC-EMMA is benchmarked by transforming the data back to the lab frame and then comparing it to data from an OSIRIS lab frame simulation. In the OSIRIS lab frame run, the laser propagates in the x_1 direction together with the moving window; as in the UPIC-EMMA run, the laser is launched from a moving

antenna. In the UPIC-EMMA simulation $\gamma_b = 14$ is used. Periodic boundary conditions are used for transverse directions in both cases. The transformed UPIC-EMMA boosted frame data (to the lab frame) are plotted together with the lab frame OSIRIS data in figure 6.5. Good agreement is found between the two cases. Note the shift in the laser wave packet between the OSIRIS data and UPIC-EMMA data. We verified that the shift was attributed to the difference in group velocity between the Yee solver and spectral solver (transformed back to lab frame).

6.3.4 Filters

Earlier the mode numbers of the fastest growing modes of the numerical Cerenkov instability in the spectral solver were identified as Eq. (6.2). Based on this equation, we use filters that eliminate a range of $\vec{k}$'s centered around this parabola. Specifically, we multiply all modes by either 1 or 0. Those modes multiplied by 0 are those in the range:

$$k_2^2 = \pm 2k_{g1}(k_1 + \Delta k_1) \tag{6.20}$$

in 2D, and

$$k_2^2 + k_3^2 = \pm 2k_{g1}(k_1 + \Delta k_1) \tag{6.21}$$

in 3D. Δk_1 is usually chosen to be

$$0.9 \times \frac{k_{g1}}{2} < |\Delta k_1| < 1.02 \times \frac{k_{g1}}{2} \tag{6.22}$$

6.4 LWFA simulations with UPIC-EMMA

We next present simulation results using UPIC-EMMA to model LWFA in a boosted frame. Two-dimensional simulations in the linear and nonlinear regimes are presented for two different choices of γ_b and the results are compared to OSIRIS simulation results in the lab frame (the UPIC-EMMA results are transformed back to the lab frame). We also present 3D results from UPIC-EMMA including comparison with OSIRIS lab frame simulations. For

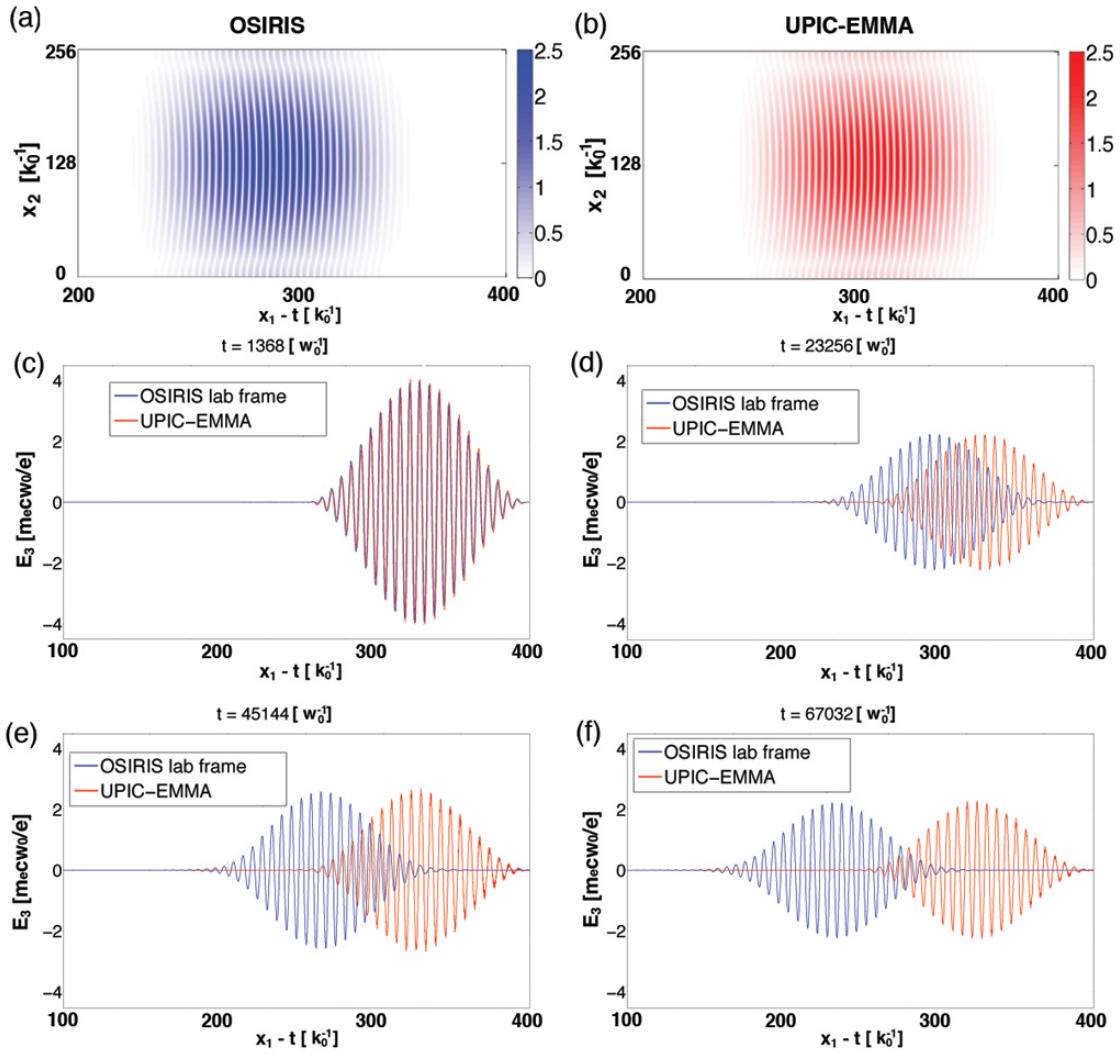

Figure 6.5: (a) is the 2D plot of the laser (polarized in x_3 direction) E_3 field at $t = 13680\,\omega_0^{-1}$, and (b) shows the laser E_3 field transformed back from the boosted frame data. (c)–(f) shows the comparison of on-axis E_3 field between OSIRIS data and UPIC-EMMA data at various time points. $x_1 - t$ is the coordinates moving together with the moving window.

the linear cases $a_0 = 0.1$ (a_0 is the normalized vector potential of the laser) while for the nonlinear cases a_0=3.0 or 4.0. Precise values for the simulations parameters are shown in tables 6.2 and 6.1.

Plasma density n_0	$8.611 \times 10^{-4} n_0 \gamma_b$
Laser	
pulse length τ	$70.64 k_0^{-1} \gamma_b (1 + \beta_b)$
pulse waist W	$117.81 k_0^{-1}$
polarization	circular
Lab frame simulation ($\gamma_b = 1$)	
grid size $(\Delta x_1, \Delta x_2, \Delta x_3)$	$(0.2 k_0^{-1}, 3.40 k_0^{-1}, 3.40 k_0^{-1})$
time step Δt	$0.199 \omega_0^{-1}$
number of grid (moving window)	$4000 \times 512 \times 512$
particle shape	quadratic
3D boosted frame simulation	
grid size $\Delta x_{1,2,3}$	$0.2 k_0^{-1} \gamma_b (1 + \beta_b)$
time step Δt	$0.04 \omega_0^{-1} \gamma_b (1 + \beta_b)$
number of grid ($\gamma_b = 17$)	$4096 \times 256 \times 256$
particle shape	quadratic

Table 6.1: Simulation parameters for the 3D simulations (related to figure 6.10). The laser frequency ω_0 and laser wave number k_0 are used to normalize simulation parameters, and $n_0 = m_e \omega_0^2 / (4\pi e^2)$.

In figure 6.6, results from the a_0=0.1 case are shown. In the top row, the wakefield E_1 is shown at various lab frame times for an OSIRIS lab frame simulations (blue) and for two UPIC-EMMA simulations where γ_b=14 (red) and γ_b=28 (green) respectively. This figure shows that the two UPIC-EMMA results agree well for all times while the OSIRIS result lines up with UPIC-EMMA results at early times and then slips backwards due to numerical dispersion. There is no evidence of noise in these plots in the wake and laser fields due to the numerical Cerenkov instability. The fact that the OSIRIS lab frame result slips backwards

is due to the numerical dispersion error in v_g that was discussed earlier. In the bottom row of figure 6.6, the laser field (E_3) is plotted at the same times. The same colors are used to show the results from the three simulations. The slippage of the OSIRIS lab frame curve is also seen in the laser field.

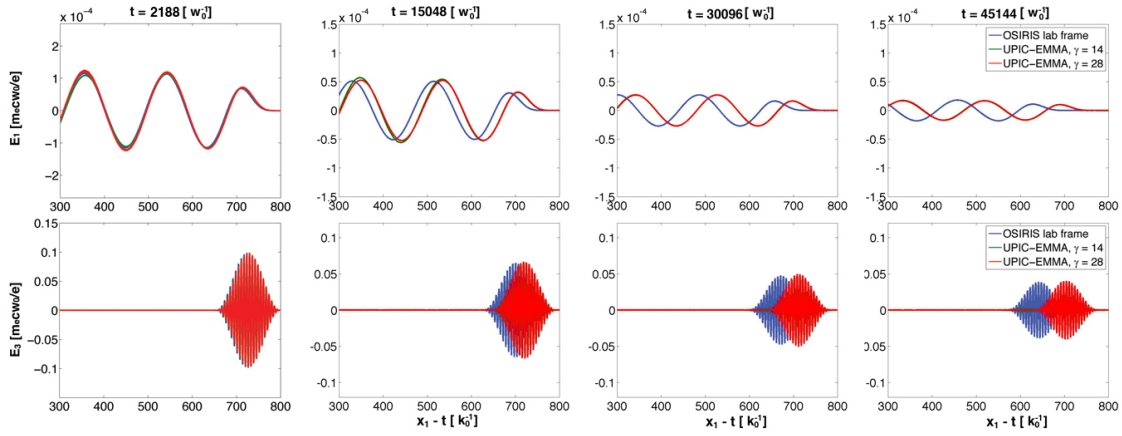

Figure 6.6: Comparison of the on-axis E_1 and E_3 between OSIRIS lab frame simulation, and UPIC-EMMA boosted frame simulation ($\gamma = 14, 28$) at various time steps, for $a_0 = 0.1$. $x_1 - t$ is the coordinates moving together with the moving window.

We next show results for a more nonlinear case where a_0=3.0. As before, there is a lab frame OSIRIS simulation and two UPIC-EMMA boosted frame simulations with γ_b=14 and γ_b=28. The same colors as in figure 6.6 are used to distinguish the data from these three simulations. We plot the accelerating field in the upper row and the laser field in the lower row at four various lab frame times (different times than used in figure 6.6). Similar to the linear $a_0 = 0.1$ case, the wakefields from the three simulations agree well at early times while for later times the OSIRIS results slip behind. While the agreement between the two boosted frame simulations is still good, it is not as good as for the previous case. The differences in the laser field are small for larger values of $x_1 - t$ (at the head of the laser) and there are differences at later times.

Next in figures 6.8 and 6.9, we present results from a case where the laser amplitude is increased to $a_0 = 4.0$. In the top row of figure 6.8, the plasma density and wakefield in the Lorentz boosted frame with $\gamma_b = 14$ are plotted at $t' = 6180\omega_0^{-1}$. There is no evidence of the

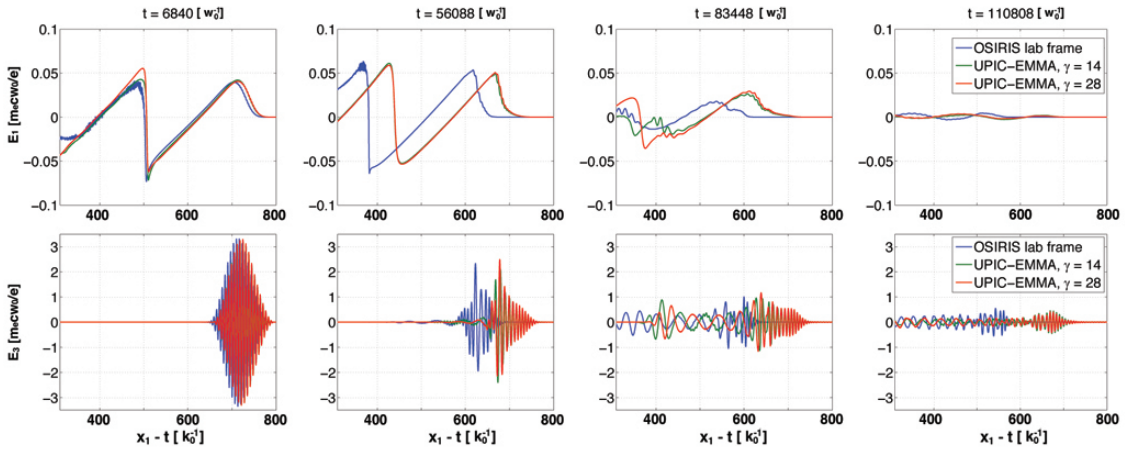

Figure 6.7: Comparison of the on-axis E_1 and E_3 between OSIRIS lab frame simulation, and UPIC-EMMA boosted frame simulation ($\gamma = 14, 28$) at various time steps, for $a_0 = 3.0$. $x_1 - t$ is the coordinates moving together with the moving window.

numerical Cerenkov instability. Only a small region of the simulation box, including where the instability is most robust, is plotted. In the lower row of figure 6.8 we also plot in the lab frame the wakefields obtained in these three simulations. As in the two previous cases, good agreement is found in the wakefield amplitude. There is slippage of the wakefield in the OSIRIS simulation and small differences between the two boosted frame simulations.

It is worth noting that the main purpose of this chapter is to illustrate the detailed simulation setups when performing LWFA Lorentz boosted frame simulation using UPIC-EMMA, which is the main content of [71]. When [71] was published, we were not fully aware of the effect of the $(\mu, \nu_1) = (0, 0)$ modes on these simulations. In fact the $a_0 = 4$ simulations discussed here were repeated in section 3.2.2, with the $(\mu, \nu_1) = (0, 0)$ eliminated by either reducing the time step, or modifying the EM dispersion relation in the spectral solver at localized $\vec{k}$ regions. By comparing Fig. 6.8 with Fig. 3.4, we can see that by eliminating the $(\mu, \nu_1) = (0, 0)$ modes, better agreements with the lab frame simulations are obtained for the self-injection.

Figure 6.9 shows the comparison of the laser E_3 fields for the three a_0=4.0 cases. In the top row we show line outs of the laser at four different propagation distances (times). The

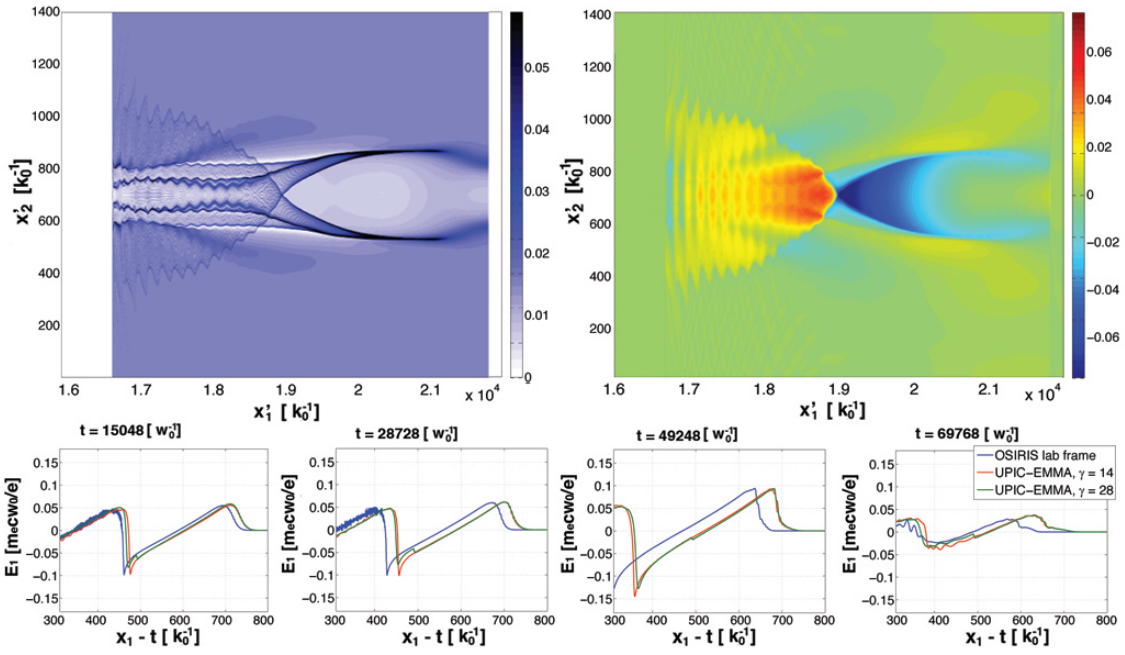

Figure 6.8: UPIC-EMMA boosted frame simulation ($\gamma = 14, 28$) for $a_0 = 4.0$. First row shows the 2D plots of plasma electron density (left), and the corresponding E_1 for $t' = 6180\ \omega_0^{-1}$ in the boosted frame ($\gamma = 14$). The second row shows the on-axis E_1 comparison between OSIRIS lab frame, and UPIC-EMMA boosted frame simulation ($\gamma = 14, 28$). $x_1 - t$ is the coordinates moving together with the moving window.

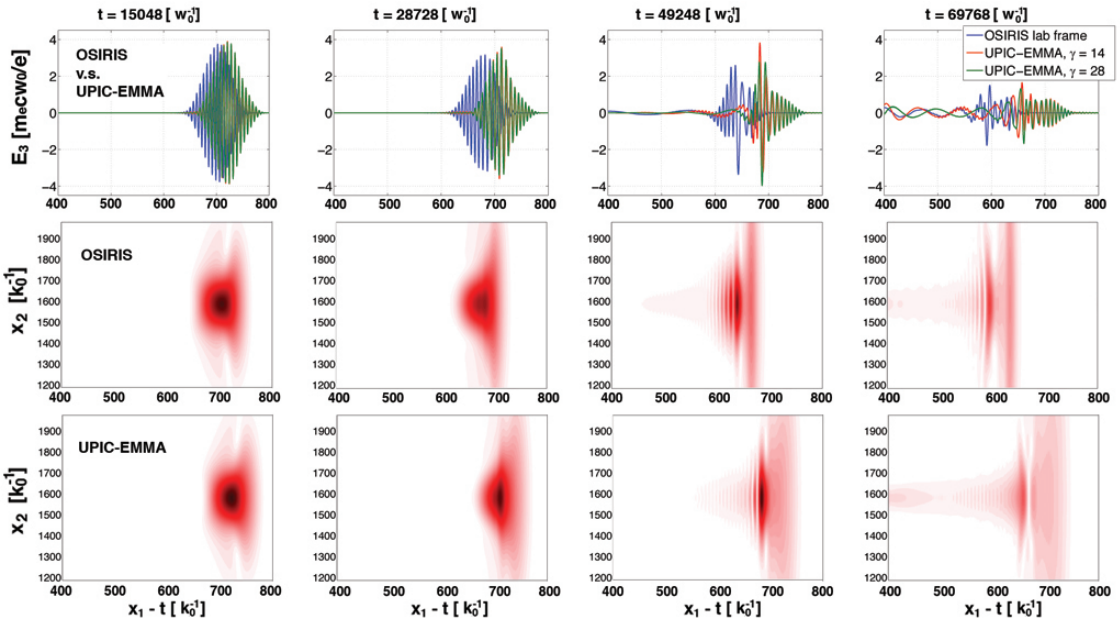

Figure 6.9: Comparison of the E_3 field between OSIRIS lab frame simulation, and UPIC-EMMA boosted frame simulation ($\gamma = 14, 28$) at various time steps, for $a_0 = 4.0$. The first row shows on-axis E_3 comparison between OSIRIS lab frame, and UPIC-EMMA boosted frame ($\gamma = 14, 28$). The second and third rows show the 2D comparison between the OSIRIS lab frame results and the transformed data from UPIC-EMMA boosted frame ($\gamma = 14$). $x_1 - t$ is the coordinates moving together with the moving window.

OSIRIS lab frame curve slips backwards. As in the other cases, the boosted frame curves line up at the front of the laser and as in the other nonlinear case differences in the curves are seen in the back of the laser. In addition, for the $\gamma_b = 14$ case we transform not only the on-axis data, but also the off-axis data in order to compare the 2D laser profile between OSIRIS lab frame run and UPIC-EMMA boosted frame run. The OSIRIS lab frame data is shown in the middle row and the UPIC-EMMA data in the bottom row. Only a part of the simulation box is shown. Good agreement is found in how the laser pump depletes between the two runs and in how the shape evolves. The slippage of the OSIRIS simulation results is seen.

Last, to illustrate that UPIC-EMMA is fully working in three-dimensions, we present the 3D results of UPC-EMMA using the simulation parameters in table 6.1 and a_0=4.0. These parameters are similar to those in Ref. [16]. In figure 6.10 (a) and (b) we present 2D slices in the boosted frame of the plasma density and wakefield at the center of the box in the $\hat{3}$-direction at $t' = 15335\omega_0^{-1}$. As in the 2D cases, no noise from the numerical Cerenkov instability is evident. In figure 6.10 (c), the wakefield at $t = 3980\ \omega_0^{-1}$ in the lab frame from the OSIRIS lab frame simulation (blue) and UPIC-EMMA boosted frame simulation (red) are shown. The curves agree well but not perfectly. Note that in (c), there is no slippage because we are showing the result at a time where little slippage has occurred. Future work will involve understanding these differences for these nonlinear cases.

6.5 Summary

In this chapter, we described the rapid development of a new 3D PIC code that can be used to model laser wakefield acceleration in the Lorentz boosted frames. In such simulations a plasma is drifting at relativistic speeds towards the laser, which leads to the numerical Cerenkov instability. The growth rates and unstable mode numbers of the numerical Cerenkov instability depends on the type of Maxwell field solver used. The new code, called UPIC-EMMA, uses a spectral field solver, and is fully parallelized. It is built using the components of the UPIC Framework, which is a set of modules for building parallelized PIC

Plasma density n_0	$1.148 \times 10^{-3} n_0 \gamma_b$
Laser	
pulse length τ	$70.64 k_0^{-1} \gamma_b (1 + \beta_b)$
pulse waist W	$117.81 k_0^{-1}$
polarization	$\hat{3}$-direction
Lab frame simulation ($\gamma_b = 1$)	
grid size $(\Delta x_1, \Delta x_2)$	$(0.2 k_0^{-1}, 2.746 k_0^{-1})$
time step Δt	$0.199 \omega_0^{-1}$
number of grid (moving window)	4000×512
particle shape	quadratic
2D boosted frame simulation	
grid size $\Delta x_{1,2}$	$0.0982 k_0^{-1} \gamma_b (1 + \beta_b)$
time step Δt	$0.0221 \omega_0^{-1} \gamma_b (1 + \beta_b)$
number of grid ($\gamma_b = 14$)	16384×512
number of grid ($\gamma_b = 28$)	8192×256
particle shape	quadratic

Table 6.2: Simulation parameters for the 2D simulations, with $a_0 = 0.1, 3.0, 4.0$ (related to figure 6.6, 6.7, 6.8, and 6.9). The laser frequency ω_0 and laser wave number k_0 are used to normalize simulation parameters, and $n_0 = m_e \omega_0^2 / (4\pi e^2)$.

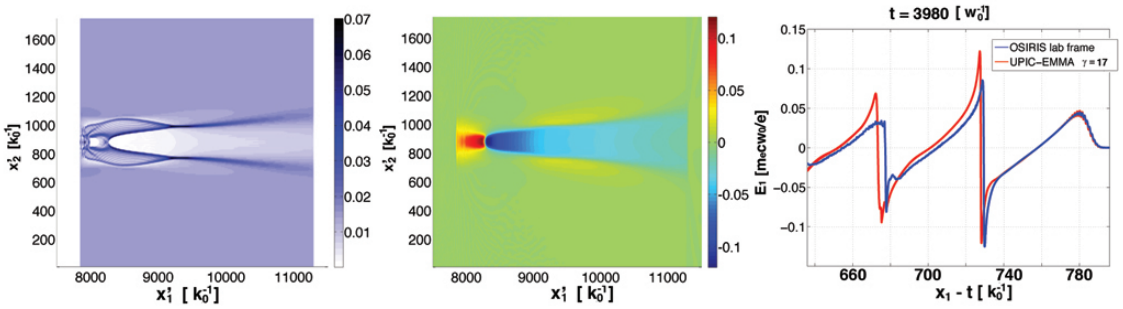

Figure 6.10: Results from 3D UPIC-EMMA boosted frame simulation ($\gamma = 17$). (a) and (b) present 2D cross section plot of the plasma electron density, and E_1 at $t' = 15335\ \omega_0^{-1}$, while (c) shows the on-axis E_1 comparison at $t = 3980\ \omega_0^{-1}$ in the lab frame. $x_1 - t$ is the coordinates moving together with the moving window.

codes with FFT based (spectral) solvers. The use of a spectral solver in which the fields are solved for in Fourier space allows for more convenient mitigation of the numerical Cerenkov instability. The phase velocity of light in vacuum and in a plasma is always greater than the speed of light for a spectral solver. In such cases, the fastest growing modes of the numerical Cerenkov instability are due to the first spatial aliased beam mode and they reside at large values of $|\vec{k}|$. These modes can be easily filtered out using a "hard" low pass or "shell" filters, thereby eliminating the fast growing modes of the instability.

We presented examples of LWFA boosted frame simulations using UPIC-EMMA. Several different values of the laser amplitude were simulated ranging from a very linear regime to a nonlinear regime. For the cases shown there was no evidence of the numerical instability and good agreement was found between OSIRIS lab frame and UPIC-EMMA boosted frame simulations. In Chapter 3 we repeated some simulations in which we eliminated the main NCI modes and differences were seen. This illustrates an important observation that even if no evidence of the NCI is seen in the field fluctuations, NCI can still alter the physics in very subtle ways. The comparison showed that the wake and laser from OSIRIS lab frame simulation slipped behind the results from the boosted frame simulations as expected from numerical dispersion errors. We showed that the dispersion errors become smaller when results are transformed back to the lab frame.

CHAPTER 7

LWFA boosted frame simulation in the quasi-3D geometry

In the previous chapter, we showed the LWFA Lorentz boosted frame simulation using UPIC-EMMA, and the comparison against the corresponding lab frame results [71]. The good agreements demonstrate the accuracy, as well as the speedups that can be provided by the Lorentz boosted frame simulation technique. As mentioned in previous chapters, recently a method to expand the fields in to azimuthal harmonics and to truncate the expansion [51, 55], called the quasi-3D algorithm, has been proposed. This can reduce the computational costs of modeling 3D problem with low azimuthal asymmetry to that on the order of 2D $r-z$ simulations. Since both the Lorentz boosted frame technique and quasi-3D algorithm provides tremendous speedups for the LWFA simulation, it would be intriguing to combine these two methods in order to combine the speedups provided by each. Similarly to full PIC simulations in the Cartesian geometry, the key to the combination is the elimination of the NCI that inevitably arises due to the unphysical coupling between Langmuir modes (main and aliasing) and EM modes of the relativistic drifting plasma in the simulations. The coupling arises in the Lorentz boosted frame between modes which are purely longitudinal (Langmuir modes) and purely transverse (EM modes) in the lab frame. The coupling occurs at specific resonances $(\omega - \mu 2\pi/\Delta t) = (k - \nu_z 2\pi/\Delta z)v_b$ where μ and ν_z are the time and space aliases and Δt and Δz are the time step and grid size respectively.

While the multi-dimensional NCI theory in Cartesian coordinates has been well described in previous chapters, there are currently no analytical expressions for the numerical dispersion relation of relativistic plasma drifting in the quasi-3D geometry. However, OSIRIS [30] simulations have shown that its behavior for the quasi-3D $r-z$ geometry is very similar

to that in Cartesian geometry. As was proposed and demonstrated in Chapter 4, a hybrid Yee-FFT solver could be used to suppress the NCI in the Cartesian and quasi-3D geometries [72]. In the regular Yee (a finite difference) solver in a quasi-3D geometry [51, 21], Maxwell equations are solved in (r, z) space for each azimuthal mode m. In the hybrid Yee-FFT solver, we perform a Fourier (discrete) transform in the drifting direction of the plasma (denote as $\hat{z}$), and solve Maxwell equations in k_z space for each mode m; meanwhile, in the $\hat{r}$ direction the derivatives are represented as second order finite difference operators on a Yee grid. The current is corrected to maintain the correctness of Gauss' Law. When Maxwell's equations are solved in this way, the corresponding NCI modes can be systematically eliminated by applying the same strategies used for a multi-dimensional spectral Maxwell solver. The fastest growing modes of the NCI at $(\mu, \nu_z) = (0, \pm 1)$ can be conveniently suppressed by applying a low-pass filter in the current, the highly localized $(\mu, \nu_z) = (0, 0)$ NCI modes can be moved away from physical modes by reducing the time step, and can be completely eliminated by modifying the EM dispersion at the k_z range where the $(\mu, \nu_z) = (0, 0)$ NCI modes are located. Furthermore, higher order spatial aliasing NCI modes can be suppressed by applying higher order particle shapes. We present OSIRIS simulation results which show that Lorentz boosted simulations of LWFA can be performed in this geometry with no evidence of the NCI. It is worth noting that recently a PIC algorithm based on a fully spectral solver in quasi-3D geometry has been proposed by Lehe et. al. [65]. This scheme was demonstrated with a single-node algorithm.

In addition, according to how the lab frame information is located in the (z', t') space, we show that the computation loads can be further reduced by applying a moving window in the boosted frame simulation. In the boosted frame the window follows the plasma as opposed to the laser, which is the case when using a moving window in the lab frame.

The remainder of this chapter is organized as follows: in section 7.1 we briefly discuss the hybrid Yee-FFT solver in quasi-3D geometry, and the corresponding NCI mitigation strategies. In section 7.2, we discuss the simulation setups for modeling LWFA in the Lorentz boosted frame. We discuss the distribution of the data needed for the reconstruction of lab frame information with an emphasis on showing that using a moving window in the direction

of the plasma drift can further reduce the computational load. We then show sample quasi-3D simulations of LWFA in the Lorentz boosted frame in section 7.3, and compare the results with the corresponding 3D boosted frame and lab frame data. In particular, we concentrate on the study of the laser driver evolution as it propagates through the plasma. Good agreement is obtained when comparing the driver evolution in lab frame against that obtained from the boosted frame simulation. This demonstrates the feasibility of combining the Lorentz boosted frame technique, the quasi-3D algorithm, and a moving window. We also discuss the preliminary speedups achieved in these sample simulations. The results are summarized in section 7.4.

7.1 Hybrid Yee-FFT solver in quasi-3D geometry

In Chapter 4 we discussed the hybrid Yee-FFT solver in quasi-3D geometry. In this solver, Maxwell equations are Fourier transformed in the drifting direction of the plasma (denoted as the $\hat{z}$ direction). The fields are solved in the corresponding (k_z, x, y) space, where conventional second order finite difference operators on a Yee mesh are used in (x, y). When Maxwell equations are solved in this way, the corresponding EM dispersion of the solver leads to NCI patterns that are very similar to those from a fully spectral Maxwell solver in which Maxwell equations are solved in multi-dimensional $\vec{k}$-space. Therefore one can systematically eliminate the NCI using approximately the same strategies developed for a fully spectral solver. Importantly, the hybrid Yee-FFT solver works for both Cartesian geometry (z, x, y), and quasi-3D geometry (z, r, ϕ).

When the Maxwell solver is modified from a standard Yee solver to a hybrid Yee-FFT solver, essentially the spatial finite difference operator in the $\hat{z}$ direction is modified from second-order accuracy (derived from its finite difference form) into a greater than N-th order accuracy. However, in OSIRIS (and most of the modern PIC codes) the $\vec{E}$ and $\vec{B}$ fields are advanced using Faraday's Law and Ampere's Law, while Gauss's Law is satisfied by applying a charge conserving current deposition scheme [23, 55, 56]. This scheme begins by calculating the current using the charge conserving current deposit scheme of [23, 56] for a purely $r - z$

code. It then uses this as a common factor in the amplitude for each azimuthal harmonic of J_z and J_r together with a factor that depends on the particle position in ϕ at the half time step; and it uses this together with the particle motion in ϕ to get J_ϕ for each harmonic (see section 3.2 of Ref. [55] for more details). If the continuity equation is rigorously satisfied at each time step then by taking the finite difference version of the divergence of Ampere's law, Gauss' Law is seen to be satisfied if it is satisfied at $t = 0$.

However, the rigorous charge conserving current deposit is known only for second order finite difference operators in the $\hat{z}$ direction. Therefore, when we use a FFT for the differential operator along $\hat{z}$ direction in Faraday's and Ampere's Law, we need to modify the current appropriately so the continuity equation is still true for the modified differential operator. To accomplish this, for each azimuthal mode of current $\vec{J}^m(z, r)$ obtained from the charge conserving current deposition scheme described in [55], we Fourier transform it along $\hat{z}$-direction, and then apply a correction with the form,

$$\tilde{J}_z^m(k_z, r) = \frac{[k]_z}{k_z} J_z^m(k_z, r) \tag{7.1}$$

where

$$[k]_z = \frac{\sin(k_z \Delta z/2)}{\Delta z/2} \tag{7.2}$$

is the second order first spatial finite difference operator. This correction ensures the satisfaction of Gauss' Law throughout the simulation, as will be discussed shortly afterwards.

Each azimuthal mode of the EM fields are initially stored in the memory in (z, r) space, and are advanced in (k_z, r) space. We Fourier transform $\vec{E}$ and $\vec{B}$ along $\hat{z}$-direction, and solve Faraday's Law and Ampere's Law for each azimuthal mode m, and each Fourier mode

k_z, using the corrected current as the source term,

$$\partial_t B_r^m = -\frac{im}{r} E_z^m - ik_z E_\phi^m \tag{7.3}$$

$$\partial_t B_\phi^m = ik_z E_r^m + \partial_r E_z^m \tag{7.4}$$

$$\partial_t B_z^m = -\frac{1}{r}\partial_r (r E_\phi^m) + \frac{im}{r} E_r^m \tag{7.5}$$

$$\partial_t E_r^m = \frac{im}{r} B_z^m + ik_z B_\phi^m - J_r^m \tag{7.6}$$

$$\partial_t E_\phi^m = -ik_z B_r^m - \partial_r B_z^m - J_\phi^m \tag{7.7}$$

$$\partial_t E_z^m = -\frac{1}{r}\partial_r (r B_\phi^m) - \frac{im}{r} B_r^m - \tilde{J}_z^m \tag{7.8}$$

Here $\vec{E}^m$, $\vec{B}^m$, and $\vec{J}^m$ are all in (k_z, r) space. Note that ∂_t and ∂_r adopt the conventional finite difference form as in the Yee solver. The code is gridless in ϕ so ∂_ϕ is replaced with im. The fields are then transformed back to (z, r) space, summed over m modes, and gathered for the particle pushing.

The reasoning behind the current correction Eq. (7.1) is that the charge conserving current deposition scheme described in [55] ensures that

$$\partial_t \rho^m(k_z, r) + i[k]_z J_z^m(k_z, r) + \nabla_r J_r^m(k_z, r) + \frac{im}{r} J_\phi^m(k_z, r) = 0 \tag{7.9}$$

where $[k]_z$ is given in Eq. (7.2) for this expression, and $\nabla_r(\,\cdot\,)$ is the second order accurate finite difference operator in $\hat{r}$ direction. Therefore this correction ensures that Gauss's Law in the hybrid solver

$$ik_z E_z^m(k_z, r) + \nabla_r E_r^m(k_z, r) + \frac{im}{r} E_\phi^m(k_z, r) = 4\pi\rho^m(k_z, r) \tag{7.10}$$

is satisfied throughout the simulation if it is satisfied at $t = 0$.

7.1.1 Elimination of Numerical Cerenkov instability

We have found previously that the NCI pattern for the quasi-3D hybrid Yee-FFT solver is similar to its counterpart in the Cartesian 3D geometry [57, 72]. As a result, we can apply approximately the same mitigation strategies used for the fully spectral solver in Cartesian geometry to systematically eliminate the NCI modes for this solver [45, 46].

We first eliminate the fastest growing $(\mu, \nu_z) = (0, \pm 1)$ modes (ν_z is the spatial aliasing in $\hat{z}$ direction) by applying a low-pass filter in the current. The filter covers the entire k_z range in which the $(\mu, \nu_z) = (0, \pm 1)$ NCI modes reside to prevent an unphysical exponential energy growth in these k_z modes. This can be efficiently accomplished since the current density is already in k_z space after the Fourier transform. For instance, in the sample simulation in section 7.3.2 we are using a low pass filter that has the following expression:

$$F(k_z) = \begin{cases} 1, & |k_z| < f_l k_{gz} \\ \sin^2\left(\dfrac{k_z - f_u k_{gz}}{f_l k_{gz} - f_u k_{gz}} \dfrac{\pi}{2}\right), & f_l k_{gz} \leq |k_z| \leq f_u k_{gz} \\ 0, & f_u k_{gz} < |k_z| \end{cases} \tag{7.11}$$

where $k_{gz} = 2\pi/\Delta z$. This filter cuts off all the k_z modes larger than $f_u k_{gz}$, while allowing modes smaller than $f_l k_{gz}$ to go through the filter. A sine square function connects the two regions to ensure a smooth filtering function.

The second fastest growing NCI modes $(\mu, \nu_z) = (0, 0)$ can be eliminated by reducing the time step, and then slightly modifying the k_z operator to create a small bump in the dispersion relation to precisely avoid intersections between the main EM modes and main Langmuir modes that are highly localized in k_z [46]. When determining the simulation time step, we first choose a time step such that the $(\mu, \nu_z) = (0, 0)$ NCI modes are significantly far away from the physical modes. The time step we used for the simulations presented in this paper are $\Delta t = \Delta z/4$. After that, we apply the $[k]_z$ modification in the highly localized $|k_z|$ range. This modification makes the growth rate of the $(\mu, \nu_z) = (0, 0)$ NCI modes to be zero in theory. The $[k]_z$modification is straightforward in a hybrid Yee-FFT solver since we are essentially solving the Maxwell equation in k_z space. In the sample simulation presented in section 7.3.2 we applied the following correction to the $[k]_z$ operator

$$[k]_z = \begin{cases} k_z + \Delta k_{\max} \cos^2\left(\dfrac{kz - k_{zm}}{k_{zl} - k_{zm}} \dfrac{\pi}{2}\right), & k_{zl} < |k_z| < k_{zu} \\ k_z, & \text{otherwise} \end{cases} \tag{7.12}$$

where $k_{zu} = 2k_{zm} - k_{zl}$, and k_{zm}, k_{zl}, $\Delta k_{\max}$ are listed in Table 7.1 and 7.2.

As for higher order NCI modes, their growth rates can be reduced if needed by applying higher order particle shapes. However, for the parameter space explored in this paper, the

higher order NCI modes have growth rates several orders of magnitude smaller than the fastest growing modes, and are not seen in the simulations even when the modes with higher growth rates are suppressed. Therefore, for the simulations presented here we used the quadratic particle shapes.

Applying the strategies described above, we can systematically mitigate the NCI modes in the quasi-3D geometry. Note the Fourier transform of the current into k_z space is not only important for the efficient filtering of the NCI modes, but also required to accurately correct (compensate) the current in k_z space to exactly match the modified Maxwell solver. It is worth noting that it is now a common practice to modify either the Maxwell solver or the field interpolation to change the EM dispersion relation in order to obtain a more desirable dispersion relation [39, 40, 42, 45, 46, 64, 72]. Within these schemes, Gauss' Law is satisfied by either directly solving it (as is the case in UPIC [20, 22, 58]), or by using a current that satisfies the continuity equation through a correction (compensation) to match the current deposition scheme with the Maxwell solver (as is the case in here and in [72]).

7.2 Simulation setups in the boosted frame

The setup of a quasi-3D LWFA simulation in a Lorentz boosted frame is almost identical to its counterpart in Cartesian 2D/3D geometry. In a boosted frame with Lorentz factor γ that moves in the propagation direction of the laser, the laser pulse is colliding with a counter-propagating relativistically drifting plasma [32, 33, 53]. Due to the Lorentz transform, the plasma density increases by γ while the total plasma column length contracts by γ. The laser wavelength and pulse length stretch by $\gamma(1+\beta)$, while its Rayleigh length contracts by γ. To avoid initializing a laser with very wide transverse size due to the contracted Rayleigh length and stretched pulse length, a moving antenna is placed at the edge of the plasma boundary to inject a laser pulse into the plasma [60, 61].

7.2.1 Relationship between lab and boosted frame data

In LWFA simulations in the lab frame (i.e., a stationary plasma) the use of a moving window [28], which only follows the physical domain near the laser, significantly reduces the computational load. The moving window essentially drops plasma sufficiently far behind the laser and adds fresh plasma in front of the laser. This is illustrated in Fig. 7.1 (a) where we plot the range of space time data from a lab frame simulation. The solid box shows the total space time area while the dashed box shows the reduced area through the use of a moving window. The moving window has a length 800 $[k_0^{-1}]$, and the simulation duration is $t_{\max} = 100000$ $[\omega_0^{-1}]$. We also show the simulation data that is dumped as colored lines. The data is dumped every 20000 $[\omega_0^{-1}]$. The red ends of the data lines indicate the starting end of the moving window, while the blue ends indicate the rear end. Connecting the red ends of the data lines, we obtain the $z - t$ relation for the head of the moving windows, $t = z$ (the speed of light c is normalized to 1). The data obtained in the lab frame (assuming the code dumps data at a constant time interval) rotates in space-time in the boosted frame since the Lorentz transform is essentially a hyperbolic rotation of coordinates in Minkowski space [33, 60]. Therefore lines of data in $\hat{z}$ taken at fixed time from a Lorentz boosted frame are rotated by the Lorentz transform, i.e, $t' = t/\gamma_b - \beta_b z'$. The slope of each data line now becomes $-\beta_b$, where $\beta_b = (1 - \gamma_b^{-2})^{-1/2}$ and each data line in the lab frame which belongs to the same point in time in lab frame is now spread over a range of t' and z'. Interestingly, when we connect the red end of each data line in the boosted frame it still has a slope of c, i.e. $t' = z'$. The range of data in the boosted frame is shown in Figs. 7.1 (b), (c) and (d). The data in Fig. 7.1 (b), (c) corresponds to $\gamma_b = 20$ while that in Fig. 7.1 (d) corresponds to $\gamma_b = 5$. In Figs. 7.1 (b) and (c) we also show the smallest area (domain enclosed by dashed lines) in t', z' space that includes the area needed to reconstruct the lab frame data for the two different values of γ_b. This illustrates that the space-time area in the boosted frame can be minimized by using a moving window in this frame. In Fig. 7.1 (b) it is seen that this window moves to the left (backwards); while in Fig. 7.1 (c) the window moves to the right (forwards). We use such moving window in the boosted frame OSIRIS simulations. Currently, UPIC-EMMA boosted frame simulations in Cartesian 2D/3D geometry uses a

stationary window [Fig. 7.1 (d)] [71].

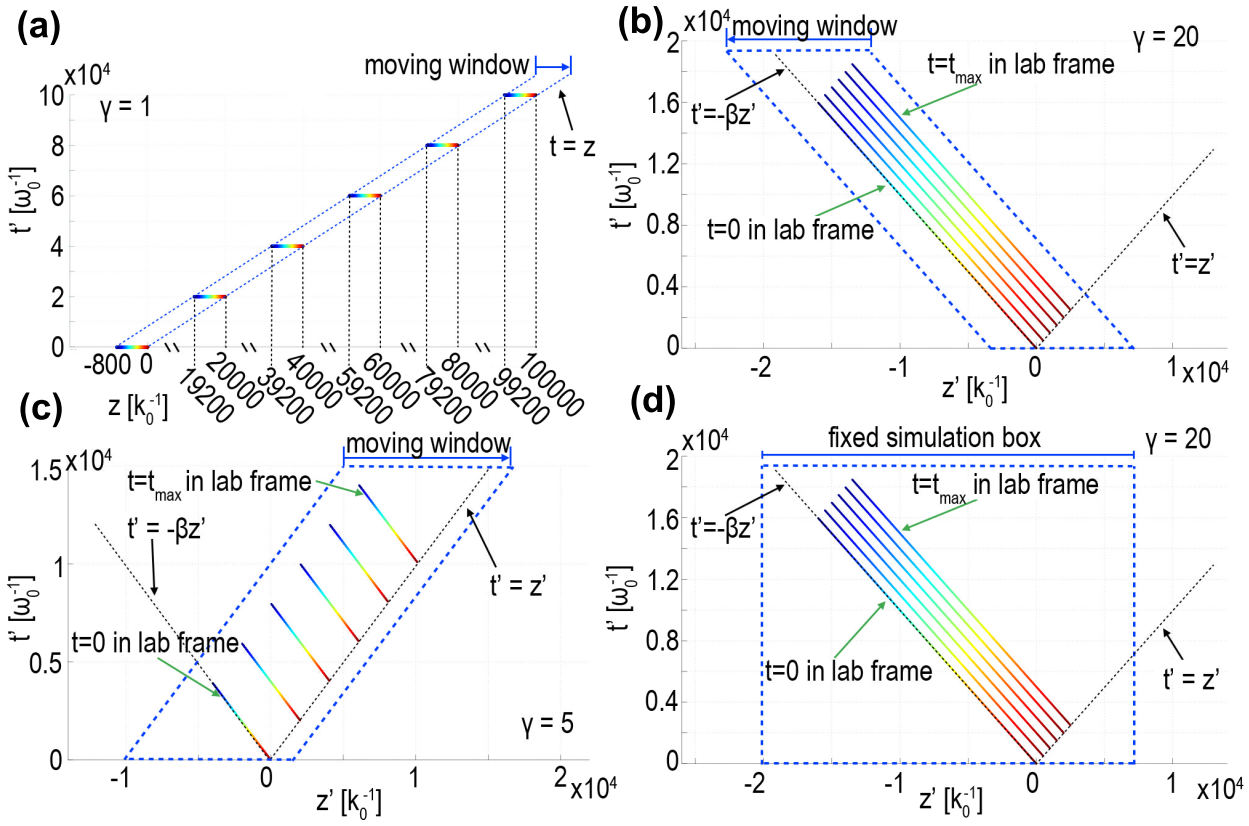

Figure 7.1: Range of important data in lab and boosted frame simulations. (a) Range of data in lab frame (stationary plasma) simulation with and without a moving window, (b) range of data in a boosted frame simulation with $\gamma_b = 20$ including with a moving window, (c) range of data in a boosted frame simulation with $\gamma_b = 5$ with a moving window, and (d) range of data in a boosted frame simulation with $\gamma_b = 20$ without a moving window.

From Fig. 7.1 it is evident that in lab frame simulations we usually dump data sparsely in time (large time intervals between time outputs), but the data at each grid is dumped at each time output. On the other hand, in order to recover the equivalent lab frame data in a boosted frame simulation, we need to sample boosted frame data at a much higher rate in time, but only need a small number of spatial locations. This can be seen by plotting a line across z' for a fixed t'. This line only intersects the equivalent lab frame data at the same number of spatial locations as the number of time outputs. We typically dump the

boosted frame data in a standard form (all grid points at small number of time steps) as well as the data needed to transform the results back to the lab frame (a small number of interpolated grid points at a large number of time steps). We then post-process the later data by performing the inverse rotation back into lab frame for comparison with the lab frame data. When running in the lab frame we also plot the necessary data needed to reconstruct the data into a boosted frame. This inverse construction method is useful during the development of a boosted frame code, as one can transform the lab frame data that has been extensively cross checked with theory, to the boosted frame, and compare the results against the results obtained by the boosted frame code.

7.2.2 Basic setup

In Fig. 7.2, we present a typical setup for a boosted frame simulation. The moving window moves from right to left following the drifting plasma. The moving antenna is also moving from right to left and injects the laser pulse from the left plasma boundary into the plasma. We place a damping section at the rear (right) end of the moving window (there is a gap between the plasma and the damping region) to damp the EM field to zero in this region. This is done because periodic boundary conditions are applied in the $\hat{z}$ direction when using the hybrid Yee-FFT solver, which requires that the EM fields need to be zero at the rear end of the simulation window to match the fields at the opposite side; otherwise the EM field at the rear end will reappear at the starting end. We note that there will be a low level of EM reflection from the damping section. In an FFT solver, the group velocity of light in vacuum is greater than the speed of light, however, since the simulation window is moving at the speed of light and the drifting plasma is drifting ultra-relativistically away from it, the reflected energy is not able to catch up with the drifting plasma. Hence the physics inside the plasma will not be affected by the reflecting EM waves. We have compared cases with the moving window plus the damping regions against cases without the moving window to confirm that the moving window plus damping region works [68]. We also note that for high γ_b boosted frame simulations, we find that the modified pusher described in Ref. [69] is required in order to get the evolution of the bubble correct. As pointed out in Ref. [69] the

usual leap frog staggering leads to issues for the Lorentz force when there is near cancellation of the electric and magnetic forces for relativistically moving particles. Determining at what γ_b the modified pusher in [69] is needed is an area of future work.

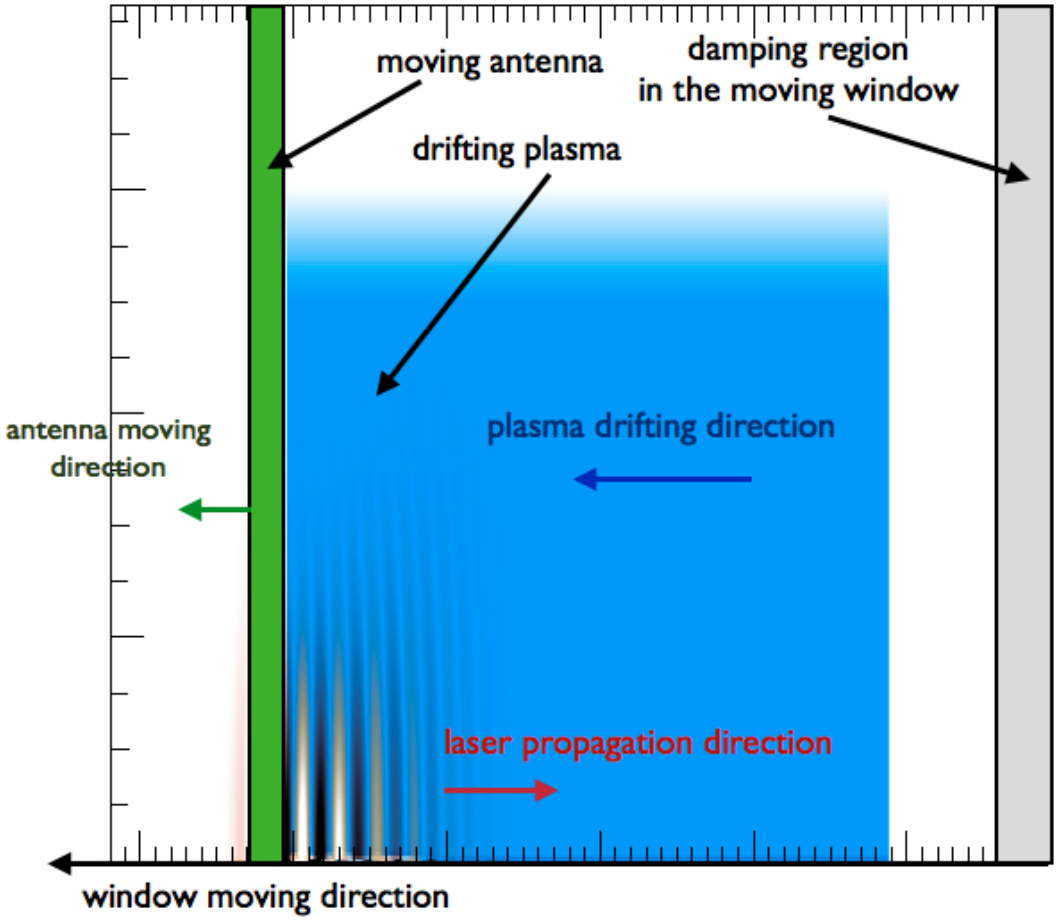

Figure 7.2: Simulation setup for a typical LWFA simulation in the boosted frame. The moving window follows the drifting plasma moving from right to left. A moving antenna injects laser pulse from left to right, and a damping region is located at the rear end of the moving window which also moves from left to right.

7.3 Sample simulations

In this section, we present two sets of sample simulations. We begin by comparing results from two boosted frame simulations where in one case we use full 3D OSIRIS and in the second case we use quasi-3D OSIRIS, in order to justify the truncation of higher azimuthal modes m in the quasi-3D boosted frame simulations. In both cases hybrid Yee-FFT solvers and the corresponding NCI mitigation schemes are used. The parameters match those in Ref. [16] whereby a 200 TW laser is focused to a spot size of 19.5 μm at the entrance of a 1.5×10^{18} cm^{-3} density plasma. The FWHM pulse length of the laser was 35 fs and the normalized vector was $a_0 = 4.0$ for a linearly polarized laser or $a_0 = 4.0/\sqrt{2}$ for a circularly polarized laser. This corresponds to a 1.3 GeV output electron energy according to the scaling laws in Ref. [16]. The numerical parameters are shown in Table 7.1. We then compare the output in the boosted frame for various azimuthal mode numbers. This comparison requires the use of a post-processing algorithm which decomposes the full 3D data into azimuthal modes [74].

We then compare the data of a LWFA boosted frame simulation in quasi-3D lab with the corresponding quasi-3D boosted frame simulation. For these simulations we explore parameters for which a full 3D lab frame simulation is not feasible due to the large CPU hours required. The parameters correspond to a 1.8 PW laser focused to a spot size of 45 μm at the entrance of a 2.5×10^{17} cm^{-3} density plasma. The FWHM pulse length of the laser was 130 fs and the normalized vector was $a_0 = 4.44$ for a linearly polarized laser or $a_0 = 4.44/\sqrt{2}$ for a circularly polarized laser. This corresponds to a 10.4 GeV output electron energy according to the scaling laws in Ref. [16]. The numerical parameters are shown in Table 7.2. The data from the boosted frame simulation is transformed back to the lab frame and it is compared against the data from the lab frame simulation.

7.3.1 3D v.s. quasi-3D boosted frame data for a 1.3 GeV case

When modeling LWFA in quasi-3D geometry, whether it is in the lab frame or boosted frame, the accelerating (E_z) and focusing fields (E_r and B_ϕ) in the bubble are mainly in

Plasma	
density n_p	$8.62 \times 10^{-4} n_0 \gamma_b$
length L	$8.0 \times 10^4 k_0^{-1}/\gamma_b$
Laser	
pulse length τ	$86.9 k_0^{-1} \gamma_b (1+\beta_b)$
pulse waist W	$153.0 k_0^{-1}$
polarization	circular
normalized vector potential a_0	4.0
Quasi-3D boosted frame simulation ($\gamma_b = 15.0$)	
grid size $\Delta z = \Delta r$	$0.1 k_0^{-1} \gamma_b (1+\beta_b)$
time step $\Delta t/\Delta x_z$	0.125
number of grid (moving window)	2048×256
particle shape	quadratic
particle per cell $(\hat{z}, \hat{r}, \hat{\phi})$	(2,2,16)
$[k]_z$ modification $([k]_{zl}, k_{zm}, \Delta k_{\max})/k_{gz}$	(0.141,0.240,0.007)
low pass filter (f_l, f_u)	(0.3,0.35)
Full 3D boosted frame simulation ($\gamma_b = 15.0$)	
grid size $\Delta z = \Delta r = \Delta y$	$0.1 k_0^{-1} \gamma_b (1+\beta_b)$
time step $\Delta t/\Delta z$	0.125
number of grid (moving window)	$2048 \times 512 \times 512$
particle shape	quadratic
particle per cell $(\hat{z}, \hat{x}, \hat{y})$	(2,2,2)
$[k]_z$ modification $([k]_{zl}, k_{zm}, \Delta k_{\max})/k_{gz}$	(0.141,0.240,0.007)
low pass filter (f_l, f_u)	(0.3,0.35)

Table 7.1: Parameters for the 3D and quasi-3D LWFA simulations in the Lorentz boosted frame (discussed in section 7.3.1). The laser frequency ω_0 and number k_0 in the lab frame are used to normalize simulation parameters. The density is normalized to the critical density in the lab frame, $n_0 = m_e \omega_0^2/(4\pi e^2)$. The normalized vector potential a_0 for the laser corresponds to linear polarization.

the $m = 0$ modes of the EM fields. On the other hand, the fields associated with the laser are associated with the $|m| = 1$ mode of fields. Therefore, by keeping at least the $|m| \leq 1$ modes the self-consistent evolution of the laser and wake fields can be examined when there is nearly azimuthal symmetry. For this comparison we truncate the azimuthal harmonics keeping only the $|m| \leq 1$ modes [51, 55]. More modes can be kept in principle to study laser hosing and asymmetric spot size effects as well as to test the convergence of the results. In addition, the results and the needed truncation can be verified by comparing LWFA boosted frame simulation results from the full 3D and quasi-3D geometries. To verify the azimuthal mode truncation, we decompose the data from the full 3D OSIRIS simulation into azimuthal harmonics and compare it against the corresponding quasi-3D simulation using the parameters listed in Table 7.1. In Fig. 7.3, we plot the azimuthal decomposition of the 3D data for E_z and E_r at $t' = 4494.99$ $[\omega_0^{-1}]$, and compare it against the corresponding quasi-3D data at the same time. For the $|m| \leq 1$ modes, very good agreement is observed. In addition, we plot the higher order $m = 2, 3$ modes from the 3D data in Fig. 7.4. We can see that the higher order modes are at least one order of magnitude smaller than those of the $m = 0, 1$ modes, which indicates that the truncation of azimuthal harmonics at $|m| \leq 1$ in the quasi-3D simulations is reasonable when the laser is nearly symmetric.

7.3.2 Quasi-3D lab frame v.s. boosted frame data for a 10.4 GeV case

Next, we compare data from a quasi-3D LWFA simulation in the lab frame against data Lorentz transformed back to the lab frame from a quasi-3D simulation. A laser with normalized vector potential of $a_0 = 4.44$ (converted to linear polarization) with pulse length of 130 fs, and spot size of 45 μm propagates into a plasma column 20.8 cm long (in the lab frame). We use a boosted frame with $\gamma_b = 26.88$, and use a moving window as described earlier that follows the relativistically drifting plasma. A moving antenna injects the laser pulse into the plasma, and a damping region absorbs the EM field at the rear end of the moving window. In the upper $\hat{r}$ boundary of the simulation box we applied the Perfectly-Matched-Layer [67] boundary condition (see Ref. [68] for more details). The plasma density is uniform along the $\hat{z}$ direction. It is uniform in the $\hat{r}$ direction from $0 \leq r \leq 7000$ $[k_0^{-1}]$ (where k_0 is the

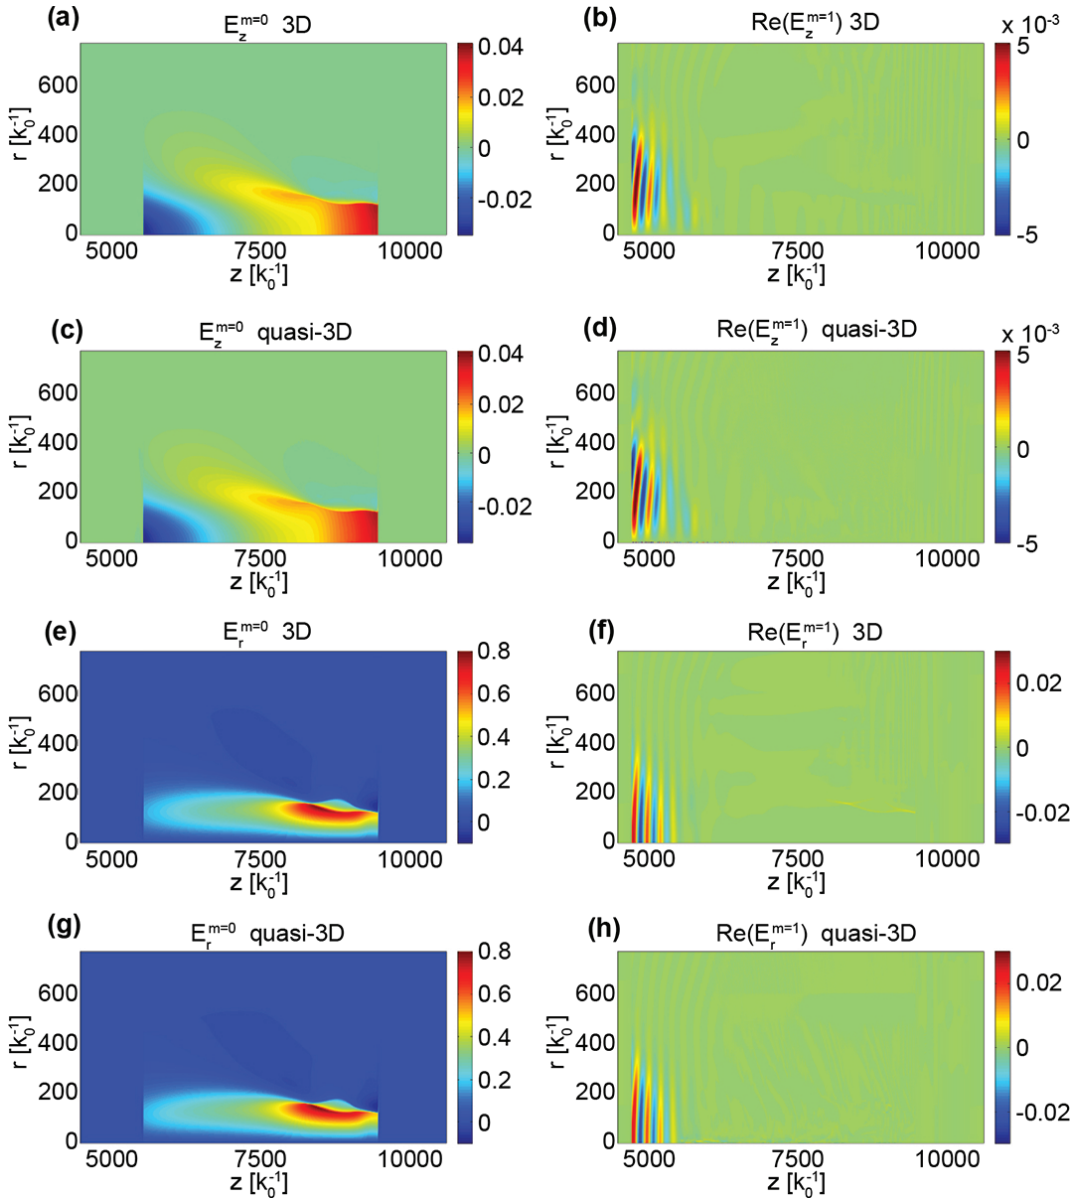

Figure 7.3: Comparison of simulation results in 3D and quasi-3D geometries for the $a_0 = 4.0$ (converted to linear polarization) 1.3 GeV LWFA stage run (as discussed in section 7.3.1). All results are from boosted frame simulations. On the left are the $m = 0$ modes of E_z and E_r. On the right are the real part of $E_z^{m=1}$ and $E_r^{m=1}$. Results from a full 3D boosted frame case are compared against a quasi-3D OSIRIS case where only $|m| \leq 1$ modes were kept. Simulation parameters are listed in Table 7.1.

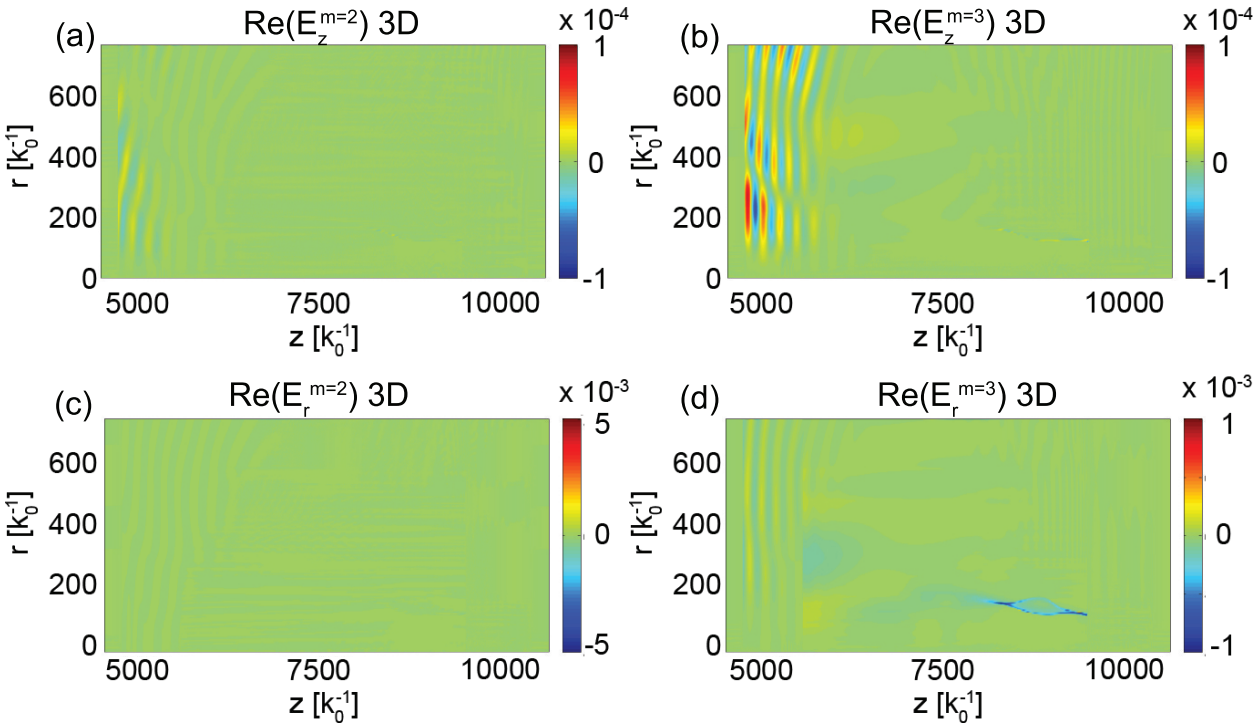

Figure 7.4: Higher order m modes of Re(E_z) and Re(E_r) obtained from a full 3D LWFA boosted frame data (as discussed in section 7.3.1). On the left are Re(E_z) and Re(E_r) for mode $m = 2$, while on the right are Re(E_z) and Re(E_r) for mode $m = 3$. The simulation parameters used are listed in Table 7.1.

wave number of the laser in the lab frame), and then the density linearly ramps to zero at $r = 8000\ [k_0^{-1}]$ near the $\hat{r}$ upper boundary (an additional gap of 500 $[k_0^{-1}]$ is left between the $\hat{r}$ upper plasma boundary and simulation box boundary). The linear plasma density ramp is used to prevent reflection when the laser moving at an oblique angle crosses the upper $\hat{r}$ plasma boundary into vacuum. Detailed simulation parameters are listed in Table 7.2.

As mentioned in section 7.2, in the boosted frame each azimuthal mode of the EM field is dumped frequently in time, and sparsely in space. The results can be transformed back to the lab frame for post-processing. In Fig. 7.5 (a) we present 2D envelope plots of the real part of the $E_r^{m=1}$ fields. The upper plot in Fig 7.5 (a) is the boosted frame simulation results (transferred back to lab frame), while the lower plot is the lab frame data. The 2D envelope of $E_z^{m=0}$ fields for the two simulations are presented in Fig 7.5 (b). In Figs. 7.5 (c)–(h) we present the corresponding line out datas from the two simulations. As we can see from Fig. 7.5 the data from the two simulations agree well with each other, except for the area around the rear of the first bubble, which indicates that the two simulations give different self-injection results. On the other hand, the laser profiles from the two cases agree extremely well [Fig. 7.5 (e)–(h)]. In Fig. 7.5 (e) line outs of $E_z^{m=0}$ at various time steps are plotted, and they show that in the transformed boosted frame data there is stronger beam loading, which indicates that more charge is self-injected into the bubble. This is likely due to the difference in statistics between the lab frame simulation and boosted frame simulation. In the boosted frame a macro-particle represents $\sim 2\gamma_b^2$ more charge than in the lab frame, while particles in the boosted frame are $\sim 2\gamma_b$ times "fatter" since the grid size in the boosted frame is $\sim 2\gamma_b$ times larger, and this could affect the self-injection process.

To confirm the differences are related to the self-injection process, we repeated the lab frame and boosted frame simulations in regimes with no self-injection, at $a_0 = 3.0$ (converted to linear polarization), while keeping the other parameters as listed in Table 7.2. In Fig. 7.6 (a) and (b) we show the line out of the wakefield $E_z^{m=0}$ at two different times in the lab frame, and in Figs. 7.6 (c) and (d) we show the corresponding amplitude envelope line outs of the laser profiles $\mathrm{Re}(E_r^{m=1})$. We see from Fig. 7.6 that for this case where there is no self-injection in the lab frame simulation, the wake field results from the lab frame

Plasma	
density n_p	$1.433 \times 10^{-4} n_0 \gamma_b$
length L	$1.63 \times 10^6 k_0^{-1}/\gamma_b$
Laser	
pulse length τ	$296.4 k_0^{-1} \gamma_b (1+\beta_b)$
pulse waist W	$351.9 k_0^{-1}$
polarization	circular
normalized vector potential a_0	4.44
Lab frame simulation ($\gamma_b = 1$)	
grid size in $(\hat{z}, \hat{r})$	$(0.2 k_0^{-1}, 4.74 k_0^{-1})$
time step $\Delta t/\Delta x_z$	0.9974
number of grid (moving window)	7920×1248
particle shape	quadratic
particle per cell $(\hat{z}, \hat{r}, \hat{\phi})$	(2,1,8)
Boosted frame simulation ($\gamma_b = 26.88$)	
grid size (square cell)	$0.2 k_0^{-1} \gamma_b (1+\beta_b)$
time step $\Delta t/\Delta x_z$	0.25
number of grid (moving window)	8192×792
particle shape	quadratic
particle per cell $(\hat{z}, \hat{r}, \hat{\phi})$	(2,2,16)
$[k]_z$ modification $([k]_{zl}, k_{zm}, \Delta k_{\max})/k_{gz}$	(0.135,0.231,0.005)
low pass filter (f_l, f_u)	(0.3,0.35)

Table 7.2: Parameters for the quasi-3D LWFA simulations in the lab frame and Lorentz boosted frame (discussed in section 7.3.2). The laser frequency ω_0 and number k_0 in the lab frame are used to normalize simulation parameters. The density is normalized to the critical density in the lab frame, $n_0 = m_e \omega_0^2/(4\pi e^2)$. The normalized vector potential a_0 for the laser corresponds to linear polarization.

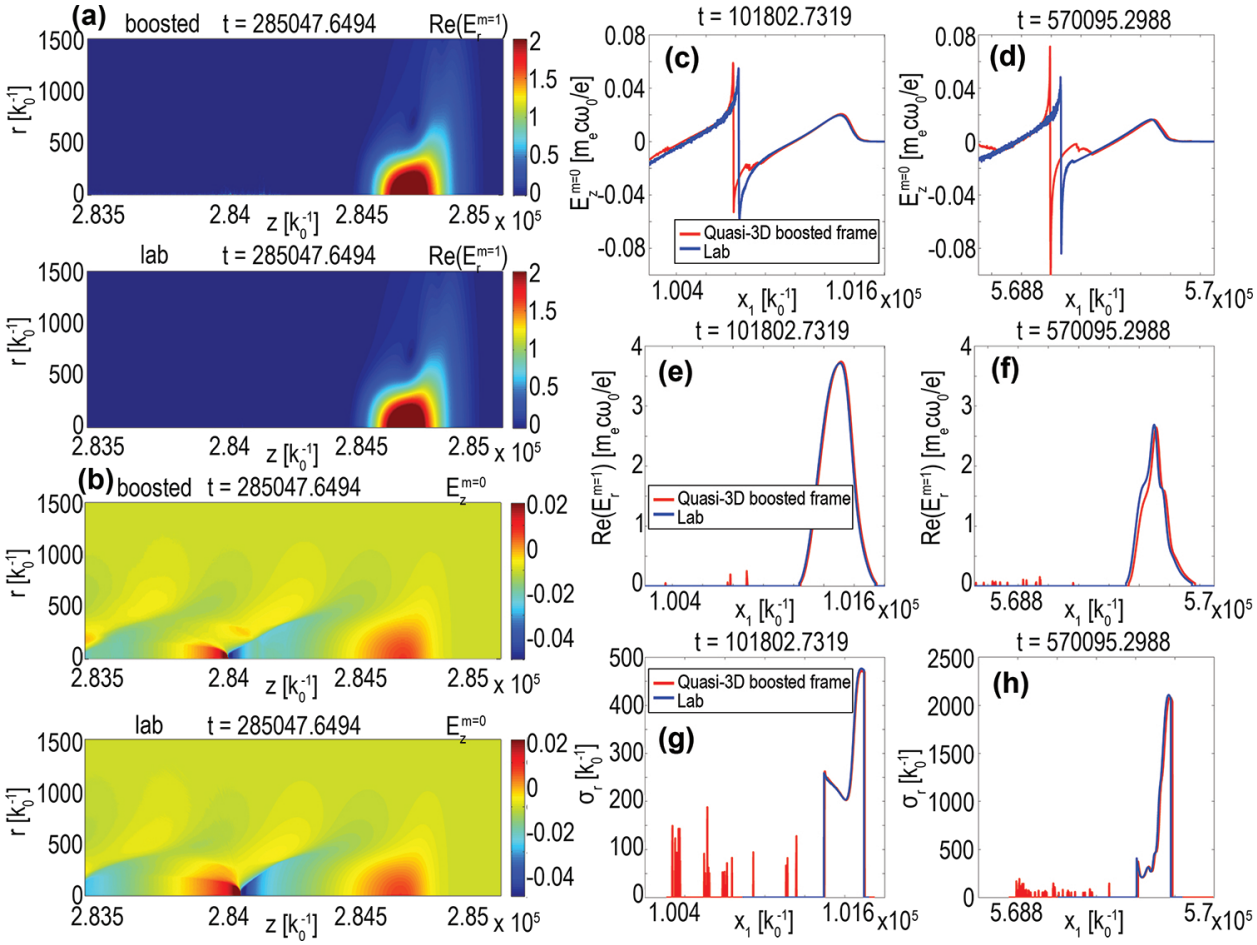

Figure 7.5: Simulation results for a $a_0 = 4.44$ (converted to linear polarization) 10.4 GeV LWFA stage run (as discussed in section 7.3.2). (a) shows the comparison of 2D envelope of Re($E_r^{m=1}$) field, which shows the evolution of laser driver as it propagates through the plasma; (b) shows the corresponding comparison of the amplitude of $E_z^{m=0}$, which shows how the wakefield of the bubble varies in the two frames due to the different self-injection results; (c), (e), and (g) are comparisons of the $E_z^{m=0}$ lineout, laser envelope Re($E_r^{m=1}$) line out, and laser spot size respectively at lab frame time $t = 101802.7\ \omega_0^{-1}$, while (d), (f), (h) are the corresponding plots at $t = 570095.3\ \omega_0^{-1}$. The simulation parameters used are listed in Table 7.2.

and boosted frame simulations agree very well. It is challenging to accurately model the self-injection process in the LWFA blowout regime. Determining the best practices for using the boosted frame technique to study self-injection at high γ_b is an area for future work.

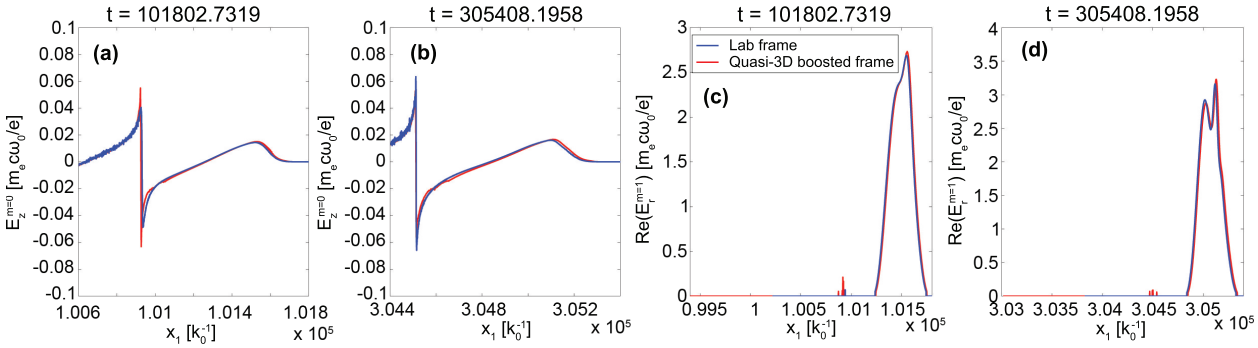

Figure 7.6: Line outs of wakefield $E_z^{m=0}$ and line outs laser field envelope of Re($E_r^{m=1}$) at various lab frame time for a $a_0 = 3.0$ case (as discussed in section 7.3.2). Since there are no self-injection in the lab frame for this case, much better agreements are obtained for the wakefield part. The simulation parameters used are listed in Table 7.2.

We plot the laser envelope and spot size obtained from the two cases in Figs. 7.5 (e)–(h). Excellent agreement can be seen for the two times presented in Fig. 7.5. Excellent agreement is also seen for the evolution of the spot size, and laser amplitude of the laser driver as it propagates through the plasma column. In Fig. 7.7 we show a detailed time history of the laser spot size and amplitude at the position of the laser where its amplitude is largest. Fig. 7.7 clearly shows that very accurate results can be obtained when using Lorentz boosted frame technique in quasi-3D geometry to study the evolution of laser driver in the plasma.

We have not yet attempted to optimize choices of parameters or the algorithm itself. However, it is still useful to compare the total CPU hours for the limited set of lab and boosted frame simulations presented in the chapter. The quasi-3D lab frame simulations presented in section 7.3.2 used around 1.6 million CPU hours. Load balancing significantly reduced the performance, and a corresponding full 3D simulation (using 8 particles per cell) would take around 300 million CPU hours in theory. Meanwhile the corresponding quasi-3D boosted frame simulation takes 2000 CPU hours. The speedup from the quasi-

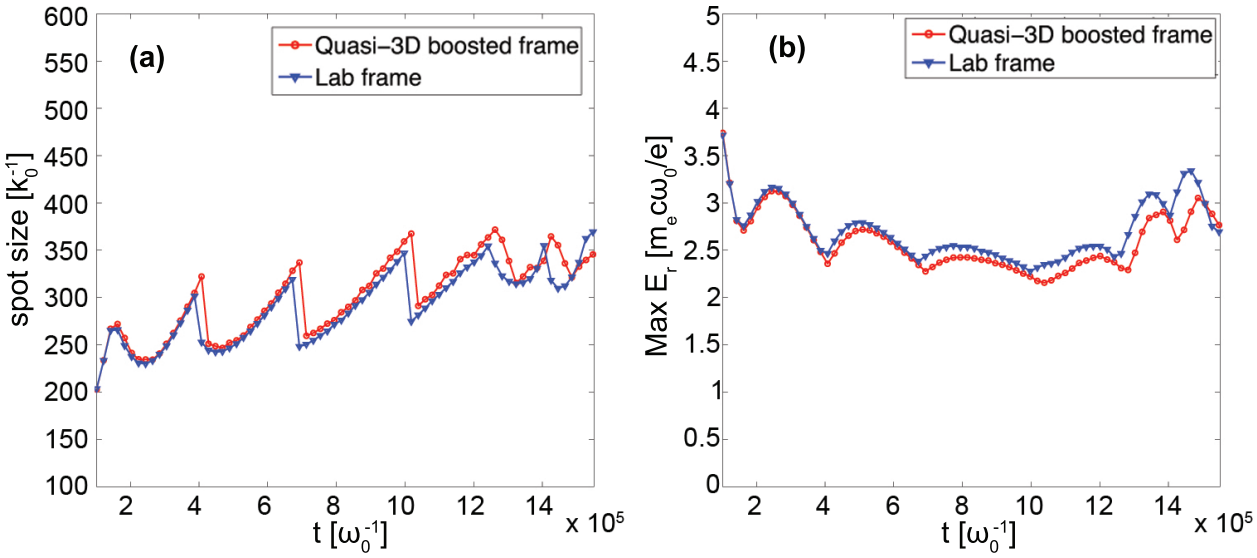

Figure 7.7: Evolution of the laser spot size and peak amplitude (discussed in section 7.3.2). (a) shows the comparison of laser spot size evolution as the laser propagates into the plasma for the two frames. The laser spot size are defined at the location where the laser has the maximum amplitude. The corresponding maximum laser amplitude evolution is shown in (b). The simulation parameters used are listed in Table 7.2.

3D lab frame is around 800. Note when calculating the speedup we take into account the fact that the transverse resolution, and particles per cell is different in the two simulations. Correspondingly the speedup achieved from the full 3D lab frame simulation to quasi-3D boosted frame is on the order of 100,000. Note if the full 3D simulation was run on a large number of processors then load balancing issues can sometimes increase the CPU hours by a factor of between 5 and 10. We note the theoretical speed up for boosted frame simulations is actually not straightforward to calculate as it depends on γ_b, the plasma length, and the laser pulse length. As noted, load imbalance is another factor that would greatly affect the speed up.

7.4 Summary

In this chapter, we described how it is possible to perform LWFA simulations in Lorentz boosted frame using the quasi-3D algorithm. The key to high fidelity Lorentz boosted frame simulations in this geometry is the use of a hybrid Yee-FFT solver that solves the Maxwell equation in k_z space in the direction that the plasma drifts, while keeping the second order finite-difference operators in the transverse directions as in a conventional Yee solver. Using this Maxwell solver we can then use the same strategies for eliminating NCI in Cartesian geometry to systematically eliminate the NCI in the quasi-3D geometry. At the same time all other features of OSIRIS are also available including single core optimization and high parallel scalability. A current correction is applied to ensure the code rigorously conserves charge. In addition, we analyzed the space-time area of the lab and boosted frame simulation data. We showed how using a moving window which follows the drifting plasma in the boosted frame the further reduce the computational load. We were able to combine Lorentz boosted frame technique with quasi-3D algorithm, together with moving window technique to achieve unprecedented speed up for the modeling of LWFA [73].

We presented comparisons of lab frame against boosted frame simulation results for a 10 GeV LWFA example that operates in the blowout regime. It was shown that the evolution of the laser driver in the plasma can be very well reproduced by the boosted frame simulation.

We also found that the self-injection process is different in the boosted frame. This is partly due to the difference in statistics between the simulations in the two frames since in the boosted frame each macro-particle represent many more real particles then in the corresponding lab frame simulation. We found excellent agreement between the lab and boosted frame results for the wake fields when a_0 was reduced to avoid self-injection.

CHAPTER 8

Dissertation summary and future work

8.1 Dissertation summary

In this dissertation, we addressed the challenge of performing PIC simulations in which a plasma (neutral or non-neutral) drifts across the grid at relativistic speeds. This is motivated by the attractiveness of performing LWFA simulations in a Lorentz boosted frame, or performing simulations of colliding plasmas in their rest frame (also a boosted frame) for studying relativistic shocks. Simulations performed in the early 1990s, and then again in the later half of the last decade revealed that a violent numerical instability occurs in such simulations. In chapter 2 we began by deriving a dispersion relation for a cold plasma drifting across the grid. We used some of the methods described in the work by Godfrey [34, 35] in the mid 1970s in which he discovered and analyzed an instability he called the Numerical Cerenkov Instability (NCI). Nonetheless, the dispersion relation in this early work [34, 35] was incomplete, and cannot be readily applied to analyze the parameter space of current interest. We derived a dispersion relation in the form of the determinant of a matrix. We then presented a way to systematically solve this equation, and obtain the pattern and growth rates of a numerical instability from the derived dispersion relation. We still refer to this instability as the numerical Cerenkov instability (NCI) as this instability can be considered as multi-dimensional version of the instability investigated in [34, 35]. We also showed that by rewriting the dispersion relation for different temporal and spatial aliasing that it was due to the unphysical coupling of modes that are purely transverse (EM) and purely longitudinal (Langmuir) in the rest frame of the drifting plasma. When using a grid these two modes are coupled by a term that vanishes in the continuous limit (infinitely small cell

size and time step).

It is found that usually the fastest growing NCI modes are the $(\mu, \nu_1) = (0, \pm 1)$ (first spatial aliasing) and $(\mu, \nu_1) = (0, 0)$ (main) NCI modes. In a Yee solver (second order accurate finite difference solver), the main NCI modes reside in a wide range of k_1 in the fundamental Brillouin zone. Moreover, they are very close to modes of physical relevance, making it difficult to eliminate them through a filter. To eliminate the main NCI modes, we found that it is advantageous to use a multi-dimensional spectral Maxwell solver, in which the electromagnetic (EM) fields are advanced in $\vec{k}$ space. For the parameter space of interest, in the spectral solver the main NCI modes become highly localized in wave number space ($\vec{k}$ space), and their growth rates are already one order of magnitude smaller than the Yee solver counterpart. Moreover, as one uses a reduced simulation time step, the main NCI modes move farther away from the $\vec{k} = \vec{0}$ (where the physical modes reside nearby), and their growth rates become even smaller. It is further found that when the main NCI modes are far away from the physical modes, one can perform a modification to the EM dispersion curve in wave number space where the NCI modes arise, thereby eliminating the intersection (coupling) between the EM modes and main Langmuir modes. Meanwhile, the $(\mu, \nu_1) = (0, \pm 1)$ modes are in a broad strip at high $|\vec{k}|$ at the edge of the fundamental Brillouin zone, and they can be eliminated through the use of a low-pass filter. The growth rates of higher order NCI couplings which are already several orders of magnitude lower than the fastest growing modes can be further reduced by using higher order particle shapes.

The scheme based on the multi-dimensional spectral Maxwell solver that effectively eliminates the NCI is both efficient, and easy to implement. However, the development of an EM-PIC code based on multi-dimensional spectral solver itself is complex. In addition, the parallel scalability of a spectral solver can be limited when the number of cells in the one direction is much larger than in other direction(s) as is often the case for the simulations of LWFA, and relativistic shock problems. This is due to the fact that a node needs global information along the direction the FFT is performed, which imposes constraints on how the simulation can be partitioned. As a result, in Chapter 4, we further explored ways to design a Maxwell solver dedicated to the elimination of the NCI. This includes a hybrid

Yee-FFT solver in which the FFT is only performed in the direction of the plasma drift. The use of an FFT in only one direction significantly reduces the complexity of implementing this scheme into an existing PIC code. However, this scheme requires care in how the current is calculated. In this algorithm (unlike the purely FFT based solver) the fields are advanced from Faraday's and Ampere's Law. Gauss' Law will also be satisfied if the current satisfied the continuity equation and Gauss' Law is satisfied at the first time step. In a finite-difference-time-domain (FDTD) solver a charge conserving current deposit is used. However, the existing algorithms conserve charge and maintain the validity of Gauss' Law for second order accurate divergence operators. If the spatial derivative is different than the second order accuracy form, the current needs to be corrected. We show that this is straightforward to do. We have shown that the NCI for such a solver is very similar to that of a pure FFT based solver. Therefore, this solver permits the use of the same strategies for effectively eliminating the NCI.

The use of an FFT in only one direction nonetheless presents challenges for parallel scalability when the FFT is performed along the direction where there are many more cells than in other direction(s). In Chapter 5, we described how an FDTD hybrid Yee-high-order solver with a customized stencil can be used to emulate an FFT based solver but with nice parallel scalability properties. This indicates that the solver can be partitioned in all directions without affecting the scalability of the code. The corresponding NCI elimination scheme includes a local 1D FFT (in x_1 direction) to the current. Transforming the current to k_1 space makes it possible to correct the current so that the Gauss' Law is satisfied for the higher order divergence operator, and to filter the current to eliminate the NCI modes at high $|k_1|$. This solver is just as effective at eliminating the NCI as the FFT-based and hybrid FFT/Finite difference solver, while preserving the scalability of the code.

We also addressed how to eliminate the NCI for a quasi-3D algorithm [51]. Since the two hybrid solvers only involve modification of the solver along the plasma drifting direction, they can be easily applied to the quasi-3D geometry. In the quasi-3D algorithm the fields and currents are defined on an $r - z$ PIC grid and expanded in azimuthal harmonics. The expansion is truncated with only a few modes so it has similar computational needs of a 2D

$r-z$ PIC code. We show that NCI has similar properties in $z-r$ as in $z-x$ slab geometry and show that the NCI can be eliminated using the hybrid FFT/second orderYee or higher order/second order Yee algorithms.

After describing how the NCI can be effectively eliminated, we apply them to both LWFA boosted frame and relativistic collisionless shock simulations. We compare results from the LWFA boosted frame to the *lab* frame. The agreement was good to excellent. With respect to LWFA boosted frame simulations, we also described how to minimize the stored data from the boosted frame in order to transform it back to the lab frame. We also proposed a method to perform LWFA boosted frame simulations by using a moving window to follow the drifting plasma (just as a moving window follows the laser as in the *lab* frame) to further reduce the computational costs. We performed LWFA Lorentz boosted frame simulations in the nonlinear blowout regime in Cartesian 2D and 3D geometry, using the FFT-based solver, and fully FDTD customized solver. In all cases we were able to perform high fidelity simulations with no evidence of the NCI. However, this comparison can be subtle as effects of the main NCI mode is not obvious. The comparison of the lab frame simulation data against the boosted frame data (transformed to the lab frame) showed excellent agreement, with tremendous speed up obtained in the boosted frame simulations. Moreover, we performed boosted frame simulations using the quasi-3D geometry and excellent agreement was found with the lab frame results. The combination of the Lorentz boosted frame simulation technique and quasi-3D algorithm via the use of hybrid solver to eliminate the NCI leads to an unprecedented speed up. There are also benefits to eliminating the NCI even for lab frame simulations, since the NCI can occur even for a relativistic electron beam propagating across the grid.

We believe the development of algorithms that effectively eliminate the NCI opens up the door for using the PIC simulation to study a wide range of physics problems that involve relativistically drifting plasma and/or relativistically drifting particle beam. Our study on the LWFA boosted frame simulation technique has made it possible to scale the electron energy from LWFA to energy orders of magnitude larger than that can be reached in the lab frame, and has made it possible to perform parameters scans for these problems.

8.2 Future work

8.2.1 NCI for $\Delta x_1 > \Delta x_{2,3}$ cells

As presented in previous chapters, there are now several schemes available to eliminate the NCI and dispersion relations exist that can accurately predict where it occurs in ω and $\vec{k}$ space for different choices in cell size, aspect ratios of the cell size, and plasma density. Much of the elimination schemes developed in this dissertation is for cases in which the cell size in the drifting direction is smaller than or equal to those in the transverse directions, i.e., $\Delta x_1 \leq \Delta x_{2,3}$ where $\hat{1}$-direction of the plasma drift. When $\Delta x_1 > \Delta x_{2,3}$, eliminating the NCI becomes more difficult. This is mainly due to the fact that the fastest growing modes, $(\mu, \nu_1) = (0, \pm 1)$ NCI modes, are no longer located at the edges of the fundamental Brillouin zone. In Fig. 8.1 we plot the location of the fastest growing modes for several aspect ratios. As we can see in Fig. 8.1, when Δx_1 is much smaller than $\Delta x_{2,3}$, the $(\mu, \nu_1) = (0, \pm 1)$ modes are spreading over the entire range of k_1, which makes it impossible to use a 1D FFT plus filtering in k_1 for the current to eliminate these modes. In this case, a straightforward modification to our elimination framework would be to apply a 2D FFT to the current, and then use a mask that pinpoints the location where the $(\mu, \nu_1) = (0, \pm 1)$ modes reside.

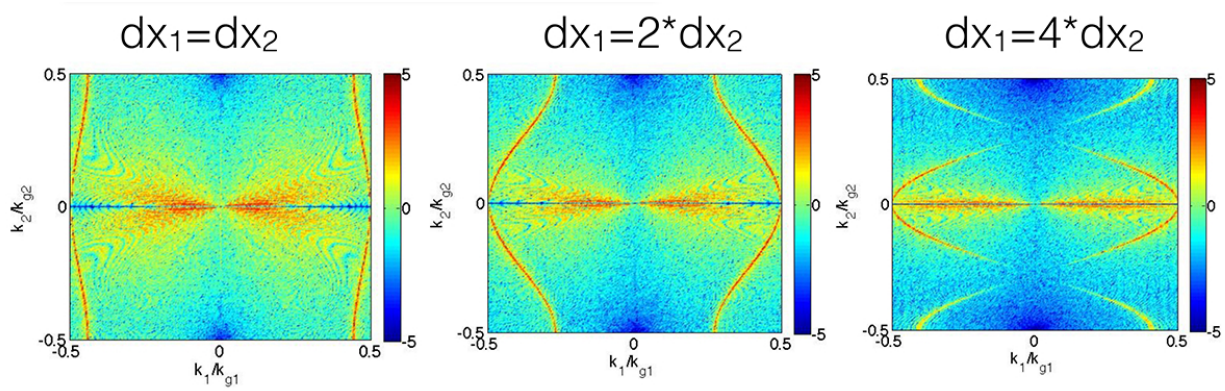

Figure 8.1: (a) shows the pattern of $(\mu, \nu_1) = (0, \pm 1)$ NCI modes for $\Delta x_1 = \Delta x_2$ in a 2D Cartesian PIC simulation of relativistic plasma drift; (b) and (c) shows the corresponding patterns for $\Delta x_1 = 2\Delta x_2$ and $\Delta x_1 = 4\Delta x_2$.

There are also quite a few important physics problems of which the PIC simulation setups would be in the $\Delta x_1 > \Delta x_{2,3}$ regime. This includes the simulation of LWFA down ramp injection in the lab frame where studying the generation of beams with ultra-high brightnesses (very narrow bunches) , and plasma wakefield acceleration (PWFA) simulations in the Lorentz boosted frame. In the former case, one needs to apply the NCI elimination scheme to mitigate the unphysical beam emittance growth due to the numerical Cerenkov radiation. Since the transverse size of the beam in the down ramp injection scheme is extremely small, one needs a much smaller transverse grid size in order to resolve the beam transverse size, and eventually one approaches the $\Delta x_1 > \Delta x_{2,3}$ limit. In the latter case, since the PWFA simulation in the lab frame utilizes square/cubic cells, when performing such simulations in the Lorentz boosted frame, the grid size in the beam driver propagation direction the grid size will be $\sim 2\gamma$ times larger, which will also reach the $\Delta x_1 > \Delta x_{2,3}$ limit. Therefore, exploring a systematic way to the eliminate the NCI for the $\Delta x_1 > \Delta x_{2,3}$ cell sizes is important for these simulation problems.

We do not doubt that additional strategies may be developed for controlling the NCI. These might include the use of local FFTs to solve the fields as well as the currents and then quantifying the errors in the longitudinal fields for different partition sizes. Furthermore, when using FFTs one can experiment with the use of different functions for $[\vec{k}]$. Another idea might be to directly deposit the current and then correct for the longitudinal part of the fields through a Boris correction. Other ideas might involve experimenting with various time advance scheme (which effectively modifies $[\omega]$ in the dispersion relation).

8.2.2 Modeling of self-injection process in the Lorentz boosted frame

As discussed in previous chapters, one of the challenges in the LWFA Lorentz boosted frame simulation is to accurately model the self-injection process. One of the most significant differences between lab frame simulation and its corresponding boosted frame simulations is that, in the latter case each macro-particle is representing much more actual plasma particles than in the lab frame simulation. Due to the density compression in the boosted frame, as

well as the stretching of the x_1 grid size, each boosted frame macro-particle represents roughly $\sim \gamma^2$ more electrons (assuming the same number of particles per cell are used). This leads to a significant difference in the statistics of self-injection in the two frames, which could lead to difference in the self-injection results. Moreover, one of the biggest unsolved physics problems in LWFA physics is the mechanism of self-injection. For one thing, it is still unknown what is the smallest physical length that the PIC simulation needs to resolve for the self-injection process, in order to accurately capture the underlying physics. Therefore, it is unknown whether the physics of the self-injection process is being under-resolved in the boosted frame simulation. We believe a systematic study of how the boosted frame simulation parameters, e.g. particle per cell, γ_b, and grid sizes, affect the self-injection process would be important to understand how well the boosted frame simulation is modeling the self-injection process.

APPENDIX A

$(-1)^{\sum_i \nu_i}$ term, interpolation tensor and finite difference operator

A.1 $(-1)^{\sum_i \nu_i}$ term due to the half-grid offset

In section 2.1 after Eq. (2.4), we mentioned that when the EM fields are staggered (such as on a Yee mesh), there is an additional $(-1)^{\sum_i \nu_i}$ term for each component $\vec{E}(\omega', \vec{k}')$ and $\vec{B}(\omega', \vec{k}')$, where $\hat{i}$ is summed over the directions for which the specific component of the EM field is staggered a half-grid offset from where charge density is defined. In this section, we explain how this additional $(-1)^{\sum_i \nu_i}$ terms occur in the dispersion relation. For simplicity, we will depict how the $(-1)^{\sum_i \nu_i}$ term is derived in 1D Cartesian geometry. The electric field $\vec{E}$ and magnetic field $\vec{B}$ are defined according to the Yee mesh, therefore in 1D Cartesian geometry, E_1, $B_{2,3}$, and j_1 have a half-grid offset along $\hat{1}$ direction, while other quantities are defined on the grid point.

We first interpolate the EM fields defined on discrete grid locations to each particle, in order to update the particle's position and momentum. The electric field $E_{1,p}$ exerted on a particle p at position x and time $t = m\Delta t$, where Δt is the time step of the simulation can be expressed as,

$$E_{1,p}(m\Delta t, x) = \sum_n E_1\left(m\Delta t, (n+\frac{1}{2})\Delta x\right) S_1\left(x - (n+\frac{1}{2})\Delta x\right) \tag{A.1}$$

where n is the grid indices along the $\hat{1}$ direction. Applying Fourier transform $\frac{1}{2\pi}\int_{-\infty}^{+\infty} dx \exp(-jkx)$

to this equation, we obtain,

$$
\begin{aligned}
& E_{1,p}(k) \\
=& \sum_n E_1\left((n+\frac{1}{2})\Delta x\right)\frac{1}{2\pi}\int_{-\infty}^{+\infty} dx S_1\left(x-(n+\frac{1}{2})\Delta x\right)\exp(-jkx) && \text{(A.2)} \\
=& \sum_n E_1\left((n+\frac{1}{2})\Delta x\right)\frac{1}{2\pi}\int_{-\infty}^{+\infty} dx S_1\left(x-(n+\frac{1}{2})\Delta x\right)\exp\left(-jk(x-(n+\frac{1}{2})\Delta x)\right) \\
& \exp\left(-jk(n+\frac{1}{2})\Delta x\right) && \text{(A.3)} \\
=& \sum_n E_1\left((n+\frac{1}{2})\Delta x\right) S_1(k)\exp\left(-jk(n+\frac{1}{2})\Delta x\right) && \text{(A.4)} \\
=& S_1(k)E_1(k) && \text{(A.5)}
\end{aligned}
$$

where

$$
S_1(k) = \frac{1}{2\pi}\int_{-\infty}^{+\infty} dx S_1\left(x-(n+\frac{1}{2})\Delta x\right)\exp\left(-jk(x-(n+\frac{1}{2})\Delta x)\right) \qquad \text{(A.6)}
$$

$$
E_1(k) = \sum_n E_1\left((n+\frac{1}{2})\Delta x\right)\exp\left(-jk(n+\frac{1}{2})\Delta x\right) \qquad \text{(A.7)}
$$

Note when substituting k with k' in Eq. (A.7), where

$$
k' = k + \nu_i k_g \qquad k_g = \frac{2\pi}{\Delta x} \qquad \nu = 0, \pm 1, \pm 2, \ldots \qquad \text{(A.8)}
$$

we obtain

$$
E_1(k') = \exp\left(-j\pi\nu\right) E_1(k) = (-1)^\nu E_1(k) \qquad \text{(A.9)}
$$

due to the E_1's half-grid offset in $\hat{1}$ direction. The $(-1)^{\sum_i \nu_i}$ terms in other quantities in multi-dimension can be obtained in a similar way.

A.2 Interpolation tensor and finite difference operator

In this appendix we will write out the explicit expressions for the interpolation tensors $\overleftrightarrow{S}$ for the fields and the currents. For a momentum conserving scheme in 3D the interpolation

tensor for the EM field after the Fourier transform can be expressed as:

$$
\begin{aligned}
&S_{E1} = s_{l,1}s_{l,2}s_{l,3}\eta_1 \quad S_{E2} = s_{l,1}s_{l,2}s_{l,3}\eta_2 \quad S_{E3} = s_{l,1}s_{l,2}s_{l,3}\eta_3 \\
&S_{B1} = \cos(\omega\Delta t/2)s_{l,1}s_{l,2}s_{l,3}\eta_2\eta_3 \quad S_{B2} = \cos(\omega\Delta t/2)s_{l,1}s_{l,2}s_{l,3}\eta_1\eta_3 \\
&S_{B3} = \cos(\omega\Delta t/2)s_{l,1}s_{l,2}s_{l,3}\eta_1\eta_2
\end{aligned} \tag{A.10}
$$

where

$$
s_{l,i} = \left(\frac{\sin(k_i\Delta x_i/2)}{k_i\Delta x_i/2}\right)^{l+1} \tag{A.11}
$$

and $\eta_i = \zeta^{\nu_i}$, $\zeta = -1$ when the EM field has a half-grid offset in the $\hat{i}$ direction, and $\zeta = 1$ when it is defined at grid point. l refers to the order ($l = 1$ is area weighting or linear interpolation for the charge). While for the energy conserving scheme, we have

$$
\begin{aligned}
&S_{E1} = s_{l-1,1}s_{l,2}s_{l,3}\eta_1 \quad S_{E2} = s_{l,1}s_{l-1,2}s_{l,3}\eta_2 \quad S_{E3} = s_{l,1}s_{l,2}s_{l-1,3}\eta_3 \\
&S_{B1} = \cos(\omega\Delta t/2)s_{l,1}s_{l-1,2}s_{l-1,3}\eta_2\eta_3 \quad S_{B2} = \cos(\omega\Delta t/2)s_{l-1,1}s_{l,2}s_{l-1,3}\eta_1\eta_3 \\
&S_{B3} = \cos(\omega\Delta t/2)s_{l-1,1}s_{l-1,2}s_{l,3}\eta_1\eta_2
\end{aligned} \tag{A.12}
$$

The space finite difference operator for the Yee solver is:

$$
[k]_i = \frac{\sin(k_i\Delta x_i/2)}{\Delta x_i/2} \tag{A.13}
$$

and is the same for electric and magnetic field.

In Karkkainen solver, the space finite difference operator for the magnetic field $[k]_{Bi}$ is the same as Eq. (A.13), while for the electric field

$$
[k]_{Ei} = c_i\frac{\sin(k_i\Delta x_i/2)}{\Delta x_i/2} \tag{A.14}
$$

where

$$
\begin{aligned}
c_1 &= \theta_1 + 2\theta_2\{\cos(k_2\Delta x_2) + \cos(k_3\Delta x_3)\} + 4\theta_3\cos(k_2\Delta x_2)\cos(k_3\Delta x_3) \\
c_2 &= \theta_1 + 2\theta_2\{\cos(k_3\Delta x_3) + \cos(k_1\Delta x_1)\} + 4\theta_3\cos(k_3\Delta x_3)\cos(k_1\Delta x_1) \\
c_3 &= \theta_1 + 2\theta_2\{\cos(k_1\Delta x_1) + \cos(k_2\Delta x_2)\} + 4\theta_3\cos(k_1\Delta x_1)\cos(k_2\Delta x_2)
\end{aligned} \tag{A.15}
$$

and

$$\theta_1 = 7/12 \qquad \theta_2 = 1/12 \qquad \theta_3 = 1/48 \tag{A.16}$$

are the tunable parameters for the Karkkainen solver [59]. The space finite difference operator for the spectral solver is

$$[k]_i = k_i \tag{A.17}$$

The time finite difference operator for the Yee, Karkkainen, and spectral solvers are the same

$$[\omega] = \frac{\sin(\omega \Delta t/2)}{\Delta t/2} \tag{A.18}$$

With respect to the current interpolation, the current interpolation tensor is approximately:

$$S_{j1} = s_{l-1,1} s_{l,2} s_{l,3} \eta_1 \quad S_{j2} = s_{l,1} s_{l-1,2} s_{l,3} \eta_2 \quad S_{j3} = s_{l,1} s_{l,2} s_{l-1,3} \eta_3 \tag{A.19}$$

We note that the expressions for $\overleftrightarrow{S_j}$ are for the charge conserving current deposition scheme of vanishing time step $\Delta t \to 0$. Nonetheless, one can obtain the corresponding $\overleftrightarrow{\epsilon}$ for the exact charge conserving scheme by replacing Eq. (2.6) with Eq. (19) and Eq. (23) of Ref. [23] and then Fourier analyzing the expression. This was done in Ref. [40]. We note that using the more exact form in the dispersion relation the instability growth rates change by only a few percent. In addition, it does not change the presence of the optimal time step, i.e, the optimal time step arises from the staggering of the fields on the mesh and not the choice of the current deposit.

Meanwhile, when the current is directly deposited (as is done in the UPIC framework), the current interpolation functions are,

$$S_{j1} = s_{l,1} s_{l,2} s_{l,3} \eta_1 \quad S_{j2} = s_{l,1} s_{l,2} s_{l,3} \eta_2 \quad S_{j3} = s_{l,1} s_{l,2} s_{l,3} \eta_3 \tag{A.20}$$

Note that in the spectral algorithm, at each time step the longitudinal electric field is directly obtained by solving Gauss' Law which effectively ensures charge conservation.

APPENDIX B

Application of NCI elimination scheme in down ramp injection

The NCI takes place not only in LWFA Lorentz boosted frame simulation, but also in lab frame simulation of plasma-based accelerator. In these lab frame simulations, the particle beams that are being accelerated inside the bubble travel in the same direction as the driver with a very high Lorentz factor γ. This will lead to numerical Cerenkov *Radiation* (NCR), and even the numerical Cerenkov *instability*, which leads to unphysical increase of the emittance and energy spread of the accelerated beams. Note this phenomenon has been previously studied in [64], in which a modified Maxwell solver is proposed to mitigate this effect.

In this section we show one simulation case that particularly suffers from the existence of NCR and NCI in the numerical system: density down ramp injection. The idea of down ramp injection is that a laser or beam driver goes through an under-dense plasma with a density down ramp [77, 78, 79]. Due to the decrease of the phase velocity of the wake, a tremendous amount of particles are injected into the bubble, and a beam with very low transverse emittance is produced [80, 81]. As a result, down ramp injection in the blowout regime is one of the most promising injection schemes when using plasma-based accelerator as the injector for FEL application. In the lab frame PIC simulation of this injection scheme, due to the relatively high γ, and the high density of the injected beam, NCR and NCI becomes a severe problem. The NCR modulates the beam energy, which leads to the unphysical growth of the beam slice energy spread. This makes the simulation results practically unusable since the slice energy spread of the beam is the most crucial parameters that determines the whether the injected beam is suitable for FEL application.

In the following, we demonstrate how NCI elimination would significantly affect the injected beam quality with a plasma wakefield acceleration (PWFA) down ramp injection simulation. A high energy beam driver propagates through a plasma density down ramp profile with a z dependence, where $\hat{z}$ is the driver propagation direction, as presented in Fig. B.1. As indicated in Fig. B.1, the plasma density profile for the down ramp injection. The plasma density decrease linearly from $n_{p,h}$ at $z = 0$ to $n_{p,l}$ at $Z = L$. $n_{p,h}$ and $n_{p,l}$, as well as other simulation parameters are listed in Table B.1.

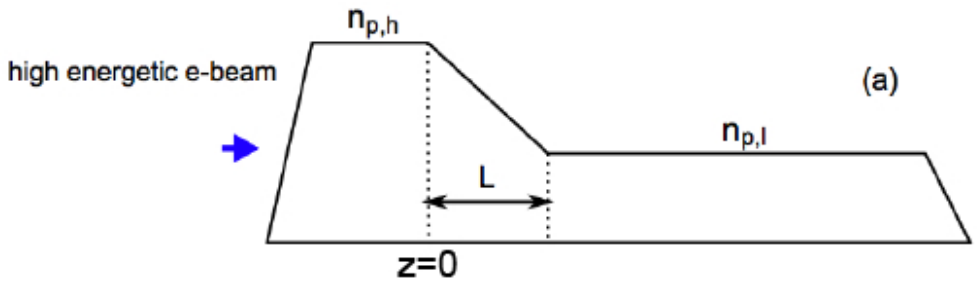

Figure B.1: The plasma density profile for the down ramp injection. The plasma density decrease linearly from $n_{p,h}$ at $z = 0$ to $n_{p,l}$ at $Z = L$.

In Fig. B.2 we present the beam energy distribution along the beam propagation direction in two 3D Cartesian OSIRIS simulations that uses the same driver and plasma parameters, yet with different simulation setups. The results in Fig. B.2 (a) uses the standard Yee solver (plus a 5-pass binary smoothing of the current), while that in Fig. B.2 (b) uses a hybrid Yee-FFT solver with the corresponding NCI elimination schemes as described in Chapter 4. As we can see in Fig. B.2 (a), the slice beam energy distribution is completely modulated by the NCR, while in Fig. B.2 (b), the NCI elimination scheme efficiently removes the unphysical modulation from the NCI, leading to a dramatically cleaner result.

Plasma density profile	
$n_{p,h}$	$1.5n_p$
$n_{p,l}$	n_p
L	$20k_p^{-1}$
Beam driver	
density	$16.0n_p$
γ	2500
RMS pulse length	$0.7k_p^{-1}$
RMS beam spot size	$0.5k_p^{-1}$
Full 3D Cartesian lab frame simulation	
grid size $\Delta z = \Delta x = \Delta y$	$0.03125k_p^{-1}$
particle shape	quadratic
particle per cell $(\hat{z}, \hat{x}, \hat{y})$	(2,2,2)
$[k]_z$ modification $([k]_{zl}, k_{zm}, \Delta k_{\max})/k_{gz}$	(0.141,0.240,0.007)
low pass filter (f_l, f_u)	(0.3,0.35)

Table B.1: Parameters for the 3D Cartesian PWFA simulation in the lab frame. Time step $\omega_p \Delta t = 0.5 k_p^{-1} \Delta z$ for the Yee solver case, and $\omega_p \Delta t = 0.25 k_p^{-1} \Delta z$ for the hybrid Yee-FFT solver case. n_p is the reference plasma density, and $\omega_p^2 = 4\pi q^2 n_p / m_e$, $k_p = \omega_p$ (c is normalized to 1).

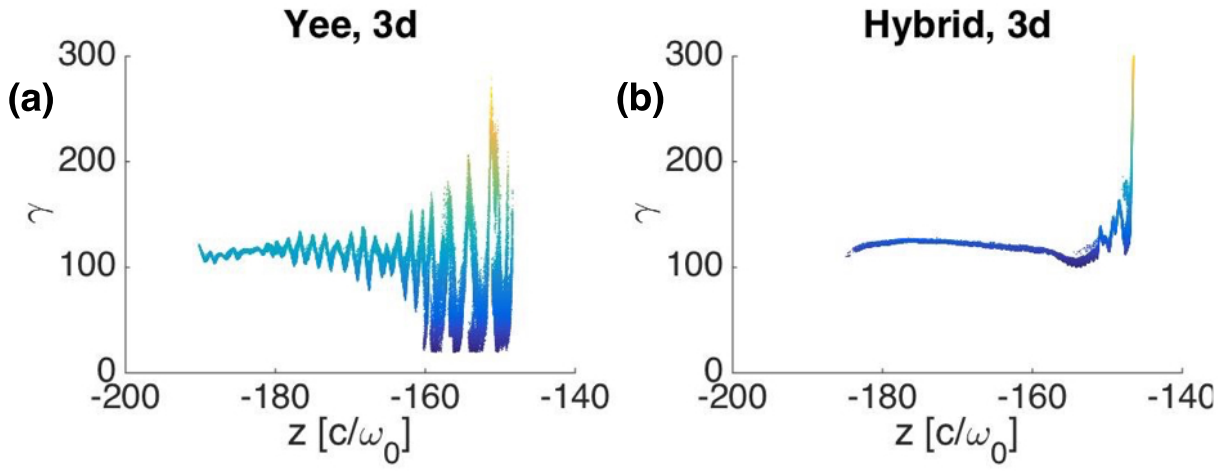

Figure B.2: The energy distribution of the injected beam along the beam propagation direction, for the two simulations. (a) uses standard Yee solver, plus a 5-pass smoothing of the current. (b) uses the hybrid Yee-FFT solver plus the NCI elimination scheme described in Chapter 4.

References

[1] O. Brüning et al., *LHC Design Report*, CERN-2004-003 (2004).

[2] A. M. Sessler, and E. Wilson, *Engines of Discovery: A Century of Particle Accelerators*, World Scientific Publishing Company, 2007.

[3] C. Pellegrini, A. M. Sessler, *The Development of Colliders (Key Papers in Physics)*, AIP Press, New York, 1995.

[4] P. Emma et al., Nature Photonics 4, 641 (2010).

[5] H. Baer et al., SLAC-R-1004 (2013).

[6] T. Tajima, J. M. Dawson, Phys. Rev. Lett. 43 (1979) 267.

[7] A. Seryi et al., in Proc. Particle Accelerator Conference 2009, Vancouver, British Columbia, Canada, 2009.

[8] E. Adli et al., SLAC-PUB-15426 (2013).

[9] S. Mangles et al. Nature 431(7008), 535 (2004).

[10] C. Geddes et al. Nature 431(7008), 538 (2004).

[11] J. Faure et al. Nature 431(7008), 541 (2004).

[12] W. P. Leemans et al., Nat. Phys. 696, 2 (2006).

[13] I. Blumenfeld el al., Nature 445, 741 (2007).

[14] M. Litos et al., Nature 92, 515 (2014).

[15] S. Corde et al., Nature 524, 442 (2015).

[16] W. Lu, el al., Phys. Rev. ST - Accel. Beams 10, 061301 (2007).

[17] A. Pukhov, J. Meyer-ter-vehn, Appl. Phys. B 74 (2002) 355.

[18] J. B. Rosenzweig, B. Breizman, T. Katsouleas, and J. J. Su, Phys. Rev. A 44, R6189 (1991).

[19] S. Lee et al., Phys. Rev. ST - Accel. Beams 5, 011001 (2002).

[20] J. M. Dawson, Rev. Modern Phys., Vol. 55, No. 2, 403 (1983).

[21] K. Yee, IEEE Trans. Antennas and Propagation, 14, 302 (1966).

[22] A. T. Lin, J. M. Dawson, H. Okuda, Phys. Fluids, 17, 1995 (1974).

[23] T. Esirkepov, Comp. Phys. Comm. 135 (2001), 144.

[24] D. F. Gordon, W. B. Mori, and T. M. Antonsen, IEEE Trans. Plasma Sci. 28, 1135 (2000).

[25] P. Mora and T. M. Antonsen, Phys. Plasma. 4, 217 (1997).

[26] C. Huang, el al., J. Comp. Phys. 217, 658 (2006).

[27] W. B. Mori, V. K. Decyk, T. C. Katsouleas, D. W. Forslund, R. D. Ferraro, E. Horowitz, *Full-Scale Numerical Experiment to Test a Future Plasma-Based Accelerator.* Submitted to NSF as part of the U.S. High Performance Computing and Communications Program (1992).

[28] C. D. Decker, W. B. Mori, Phys. Rev. Lett., Vol. 72, 490 (1994).

[29] J. -L. Vay, Phys. Rev. Lett. 98, 130405 (2007).

[30] R. A. Fonseca, et al., in: P.M.A. Sloot, et al. (Eds.), ICCS, in: Lect. Notes Comput. Sci., Vol. 2331, 2002, pp. 342–351; R. Hemker, Ph.D. thesis, UCLA, 2000 (arXiv:1503.00276).

[31] J.-L. Vay, D.P. Grote, R.H. Cohen, A. Friedman, Comp. Sci. & Discovery 5 (1) (2012), p. 014019.

[32] S. F. Martins, R. A. Fonseca, W. Lu, W. B. Mori and L. O. Silva, Nat. Phys. 6, 311 (2010).

[33] S. F. Martins, R. A. Fonseca, L. O. Silva, W. Lu, W. B. Mori, Comp. Phys. Comm. 181, 869 (2010).

[34] B. B. Godfrey, J. Comp. Phys. 15, 504 (1974).

[35] B. B. Godfrey, J. Comp. Phys. 19, 58 (1975).

[36] W. Mori and T. Katsouleas, private communications.

[37] C. K. Birdsall, A. B. Langdon, Plasma Physics via Computer Simulation, McGraw Hill, New York, 1985.

[38] E. L. Lindman, J. Comp. Phys. 5, 13 (1970); A. B. Langdon, J. Comp. Phys. 6, 247 (1970).

[39] P. Yu et.al, in Proc. 15th Advanced Accelerator Concepts Workshop, Austin, TX, 2012, in AIP Conf. Proc. 1507, 416 (2012).California, 2014.

[40] B. B. Godfrey and J.-L. Vay, J. Comp. Phys., 248, 33 (2013).

[41] Brendan B. Godfrey, J.-L. Vay, I. Haber, J. Comp. Phys., 258, 689 (2014).

[42] Brendan B. Godfrey, J.-L. Vay, J. Comp. Phys., 267, 1 (2014).

[43] B. B. Godfrey, J.-L. Vay, and I. Haber, IEEE Trans. Plas. Sci., 42 (2014), p. 1339.

[44] Brendan B. Godfrey, J.-L. Vay, Comp. Phys. Comm., 196, 221 (2015).

[45] X. Xu, el al., Comp. Phys. Comm., 184, 2503–2514 (2013).

[46] P. Yu, el al., Comp. Phys. Comm., 192, 32 (2015)052703 (2014).

[47] S. F. Martins, R. A. Fonseca, W. B. Mori, L. O. Silva, Astrophys. J. Lett. 695, L189–L193 (2009).

[48] F. Fiuza, R. A. Fonseca, J. Tonge, W. B. Mori, and L. O. Silva, Phys. Rev. Lett., 108, 235004 (2012).

[49] A. Spitkovsky, Astrophys. J. Lett. 682, L5 (2008).

[50] L. Sironi, A. Spitkovsky, J. Arons, Astrophys. J., 771, 54 (2013).

[51] A. Lifschitz, el al., J. Comp. Phys. 228 (2009), 1803.

[52] J. Vieira, private communication.

[53] J. -L. Vay, C. G. R. Geddes, E. Cormier-Michel, D. P. Grote, J. Comp. Phys. 230, 5908 (2011).

[54] J.-L. Vay, I. Haber, B. B. Godfrey, J. Comp. Phys. 243 (2013), 260.

[55] A. Davidson, el al., J. Comp. Phys. 281(2015), 1063.

[56] J. Villasenor, and O. Buneman, Comp. Phys. Comm. 69 (1992) 306.

[57] P. Yu, el al., in Proc. 16th Advanced Accelerator Concepts Workshop, San Jose, CA.

[58] V. K. Decyk, Comp. Phys. Comm. 177, 95 (2007).

[59] M. Karkkainen, E. Gjonaj, T. Lau, T. Weiland, in Proc. International Computational Accelerator Physics Conference, Chamonix, France, 2006, pp. 35–40.

[60] J.-L. Vay, C. G. R. Geddes, E. Esarey, C. B. Schroeder, W. P. Leemans, Phys. Plasmas 18, 123103 (2011).

[61] J.-L. Vay, el al., in Proc. 14th Advanced Accelerator Concepts Workshop, Annapolis, MD, AIP Conference Proceedings, 244-9 (2010).

[62] J.-L. Vay, C. G. R. Geddes, E. Cormier-Michel, Phys. Plasmas 18, 030701 (2011).

[63] C. Benedetti, et al., Nucl. Inst. and Methods in Phys. Res. A 608, 94 (2009).

[64] R. Lehe, el al., Phys. Rev. ST - Accel. Beams 16, 021301 (2013).

[65] R. Lehe, el al., Comp. Phys. Comm. 203, 66 (2016).

[66] E. S. Weibel, Phys. Rev. Lett., 2, 83 (1959); B. D. Fried, Phys. Fluids 2, 337 (1959); A. Spitkovsky, Astrophys. J. 673, L39 (2008).

[67] J. -P. Berenger, J. Comp. Phys. 114, 185 (1994).

[68] P. Yu, el al., in Proc. International Particle Accelerator Conference 2015, Richmond, Virginia, USA.

[69] J.-L. Vay, Phys. Plasmas 15, 056701 (2008).

[70] C. H. Lai, T. Katsouleas, W. B. Mori, D. W. Whittum, IEEE Trans. Plasma Phys., Vol 21, 1–8 (1993).

[71] P. Yu, el al., J. Comp. Phys. 266, 124 (2014).

[72] P. Yu, el al., Comp. Phys. Comm., 197, 144 (2015).

[73] P. Yu, el al., J. Comp. Phys. 316, 747 (2016).

[74] T. Dalichaouch, el al., Bulletin of the American Physical Society 60 (2015).

[75] I. R. Khan and R. Ohba, J. Comput. Appl. Math., 107, 179 (1999).

[76] I. R. Khan, R. Ohba, and N. Hozumi, J. Comput. Appl. Math.,150, 303 (2003).

[77] T. Katsouleas, Phys. Rev. A 33, 2056 (1986).

[78] S. Bulanov, N. Naumova, F. Pegoraro, and J. Sakai, Phys. Rev. E 58, R5257 (1998).

[79] H. Suk et al., Phys. Rev. Lett. 86, 1011 (2001).

[80] X. Xu et al., in preparation.

[81] J. Grebenyuk, et al., Nucl. Instrum. Methods Phys. Res. Sect. A 740 (2014) 246.