LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-429072

DOD-SBIR Structured Multi-Resolution PIC Code for Electromagnetic Plasma Simulations, Final Report

J. L. Vay, D. P. Grote, A. Friedman

April 26, 2010

Disclaimer

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

DOD-SBIR "Structured Multi-Resolution PIC Code for Electromagnetic Plasma Simulations"†

Final report

J.-L. Vay[1,3], D. P. Grote[2,3], A. Friedman[2,3]

[1]Lawrence Berkeley National Laboratory, CA, USA
[2]Lawrence Livermore National Laboratory, CA, USA
[3]Heavy Ion Fusion Sciences Virtual National Laboratory

April 16, 2010

ABSTRACT

A novel electromagnetic solver with mesh refinement capability was implemented in Warp. The solver allows for calculations in 2-1/2 and 3 dimensions, includes the standard Yee stencil, and the Cole-Karkkainen stencil for lower numerical dispersion along the principal axes. Warp implementation of the Cole-Karkkainen stencil includes an extension to perfectly matched layers (PML) for absorption of waves, and is preserving the conservation property of charge conserving current deposition schemes, like the Buneman-Villanesor and Esirkepov methods. Warp's mesh refinement framework (originally developed for electrostatic calculations) was augmented to allow for electromagnetic capability, following the methodology presented in [1] extended to an arbitrary number of refinement levels. Other developments include a generalized particle injection method, internal conductors using stair-cased approximation, and subcycling of particle pushing. The solver runs in parallel using MPI message passing, with a choice at runtime of 1D, 2D and 3D domain decomposition, and is shown to scale linearly on a test problem up-to 32,768 CPUs. The novel solver was tested on the modeling of filamentation instability, fast ignition, ion beam induced plasma wake, and laser plasma acceleration.

† **This work performed under the auspices of the U.S. Department of Energy by Lawrence Berkeley National Laboratory under Contract DE-AC02-05CH11231 and by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.**

I - Electromagnetic Solver Developments

New 3D and 2-1/2D electromagnetic solvers were implemented in Warp, using the standard Yee scheme or a novel stencil with lower numerical dispersion. Boundary conditions are reflective, periodic or open (PML [2]). It is parallelized using MPI message passing with1, 2 or 3-D domain decomposition.

The solver is modular by design, making use of FORTRAN derived types, simplifying by construction the implementation of adaptive mesh refinement. The main component of the structure is a 'block' containing 1 'core', 6 'sides', 12 'edges' and 8 'corners'. The core, the sides, the edges and the corners are pointers of type 'field', which encapsulate the field arrays and the field type (at present, field types are standard Yee or PML Yee). Using this structure, the implementation of boundary conditions such as Perfectly Matched Layers (PML) or periodic, is performed by either allocating each field pointer appropriately for open boundaries, or pointing them to the core object for periodic boundary conditions.

I-1 - Implementation of a "low-dispersion" solver

In [3,4], Cole introduced an implementation of the source-free Maxwell's wave equations for narrow-band applications based on non-standard finite-differences (NSFD). In [5], Karkkainen et al adapted it for wideband applications. At the Courant limit for the time step and for a given set of parameters, the stencil proposed in [5] has no numerical dispersion along the principal axes, provided that the cell size is the same along each dimension (i.e. cubic cells in 3D). The solver from [5] was modified to be consistent with the Particle-In-Cell methodology and implemented in the code Warp.

The "Cole-Karkkainen" (CK) solver [5] uses a non-standard finite difference formulation (extended stencil) of the Maxwell-Ampere equation, which reads for the electric field component E_x (from [5]):

$$
\begin{aligned}
E_x|^{n+1}_{i+1/2,j,k} = E_x|^{n}_{i+1/2,j,k} \quad & - \\
\alpha\frac{\Delta t}{\epsilon_0} D_{z,0} H_y|^{n+1/2}_{i+1/2,j,k} \quad & - \\
4\beta\frac{\Delta t}{\epsilon_0} D_{z,1} H_y|_{i+1/2,j,k} \quad & - \\
4\gamma\frac{\Delta t}{\epsilon_0} D_{z,2} H_y|^{n+1/2}_{i+1/2,j,k} \quad & + \\
\alpha\frac{\Delta t}{\epsilon_0} D_{y,0} H_z|^{n+1/2}_{i+1/2,j,k} \quad & + \\
4\beta\frac{\Delta t}{\epsilon_0} D_{y,1} H_z|^{n+1/2}_{i+1/2,j,k} \quad & + \\
4\gamma\frac{\Delta t}{\epsilon_0} D_{y,2} H_z|^{n+1/2}_{i+1/2,j,k}. &
\end{aligned}
$$

where $D_{z,0}$, $D_{z,1}$ and $D_{z,2}$ are the difference operators defined as

$$
\begin{aligned}
D_{z,0}H_y|^{n+1/2}_{i+1/2,j,k} &= \\
\frac{H_y|^{n+1/2}_{i+1/2,j,k+1/2} - H_y|^{n+1/2}_{i+1/2,j,k-1/2}}{\Delta z} &, \\
D_{z,1}H_y|^{n+1/2}_{i+1/2,j,k} &= \\
\frac{H_y|^{n+1/2}_{i+1/2,j+1,k+1/2} - H_y|^{n+1/2}_{i+1/2,j+1,k-1/2}}{4\Delta z} &+ \\
\frac{H_y|^{n+1/2}_{i+1/2,j-1,k+1/2} - H_y|^{n+1/2}_{i+1/2,j-1,k-1/2}}{4\Delta z} &+ \\
\frac{H_y|^{n+1/2}_{i+3/2,j,k+1/2} - H_y|^{n+1/2}_{i+3/2,j,k-1/2}}{4\Delta z} &+ \\
\frac{H_y|^{n+1/2}_{i-1/2,j,k+1/2} - H_y|^{n+1/2}_{i-1/2,j,k-1/2}}{4\Delta z} &, \\
D_{z,2}H_y|^{n+1/2}_{i+1/2,j,k} &= \\
\frac{H_y|^{n+1/2}_{i+3/2,j+1,k+1/2} - H_y|^{n+1/2}_{i+3/2,j+1,k-1/2}}{4\Delta z} &+ \\
\frac{H_y|^{n+1/2}_{i-1/2,j-1,k+1/2} - H_y|^{n+1/2}_{i-1/2,j-1,k-1/2}}{4\Delta z} &+ \\
\frac{H_y|^{n+1/2}_{i+3/2,j-1,k+1/2} - H_y|^{n+1/2}_{i+3/2,j-1,k-1/2}}{4\Delta z} &+ \\
\frac{H_y|^{n+1/2}_{i-1/2,j+1,k+1/2} - H_y|^{n+1/2}_{i-1/2,j+1,k-1/2}}{4\Delta z} &.
\end{aligned}
$$

and $D_{y,0}$, $D_{y,1}$ and $D_{y,2}$ are obtained from $D_{z,0}$, $D_{z,1}$ and $D_{z,2}$ by permuting the indices. The numerical dispersion along the principal axis and diagonals for cubic cells, at the Courant limit, is shown in Fig. 1 for the Yee and the Cole-Karkkainen solvers. At the Courant limit, the Yee and the CK solver offer no numerical dispersion along the 3D diagonal and the principal axes respectively. The CK solver offers overall a lower level of numerical dispersion. It is expected that this will reduce the amplitude of errors due to mismatch of numerical dispersion between grids of different resolutions when using mesh refinement.

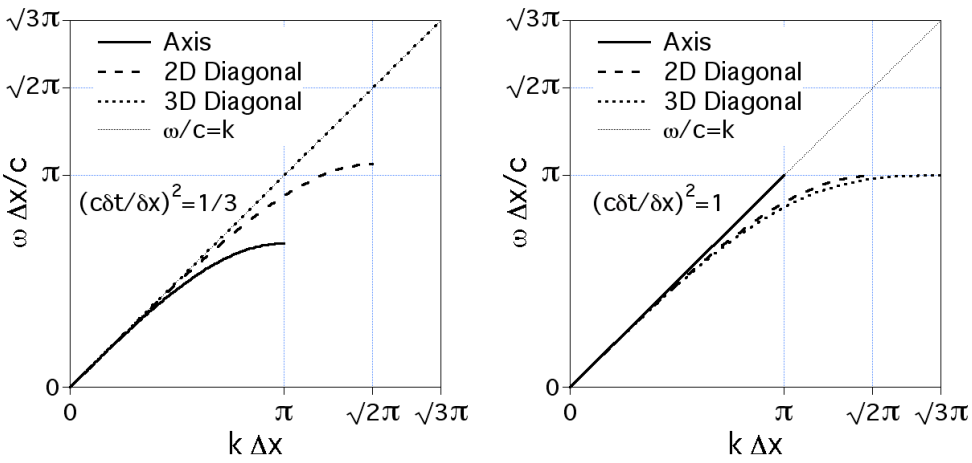

Fig. 1: numerical dispersion along the principal axis and diagonals for cubic cells at the Courant limit for the Yee and the Cole-Karkkainen solver.

For implementation into a Particle-In-Cell code, the formulation must introduce the source term into Cole-Karkkainen's source free formulation in a consistent manner. However, modifying the NSFD formulation of the Maxwell-Ampere equation so that it includes the source term in a way that is consistent with the current deposition scheme is challenging. To circumvent this problem, Warp implementation departs from Karkkainen's by applying the enlarged stencil on the Maxwell-Faraday equations, which does not contain any source term but is formally equivalent to the source-free Maxwell-Ampere equation.

In most applications, it is essential to prevent accumulations of errors to the discretized Gauss' Law. This is accomplished by providing a method for depositing the current from the particles to the grid which is compatible with the discretized Gauss' Law, or by providing a mechanism for "divergence cleaning" [7-9]. For the former, schemes which allow a deposition of the current that is exact when combined with the Yee solver is given in [10] for linear form factors and in [11] for higher order form factors. Since the discretized Gauss' Law and Maxwell-Faraday equation are the same in our implementation as in the Yee solver, charge conservation is readily verified using the current deposition procedures from [10] and [11], and this was verified numerically.

The ability is also given to the user of setting the solver adjustable coefficients, providing tunability of the numerical properties of the solver to fit the requirements of a particular application. A more detailed description of the solver and its properties is given in [12].

I-2 - Implementation of a Perfectly-Matched-Layer

I-2-a – Split PML

The split form of Perfectly Matched Layer (PML) [2] formulation applies readily to the Cole-Karkkainen form of the Maxwell equations, and was implemented in Warp with

the Yee and the CK stencil. The PML using the CK stencil was tested and compared to the standard Yee implementation in 2D and 3D. Fig. 2 shows snapshots from 2D simulations of the reflected residue from a PML layer of a pulse with amplitude given by the Harris pulse $H(t)=(10-5\cos(2\pi ct/L)+6\cos(4\pi ct/L)-\cos(6\pi ct/L))/32$, where t is time, c is the speed of light and $L=50\delta x$ is the pulse length in cell size units. A grid of 400x400 cells was used with $\delta x=\delta y$. The absorbing layer was 8 cells deep and the dependency of the PML coefficients with the index position i in the layer was $\sigma_i= \sigma_m\ (i\ \delta x/\Delta)^n$ with $\sigma_m=4/\delta x$, $\Delta=5\delta x$ and n=2. The alternative prescription for the coefficients given in [13,14], which reads $\sigma_i^*=(\xi_{i+1/2}-1/\xi_i)/\delta x$ with $\xi_i =\exp(-\sigma_i\ \delta t)$ and $\sigma_i= \sigma_m\ (i\ \delta x/\Delta)^n$, was also tested.

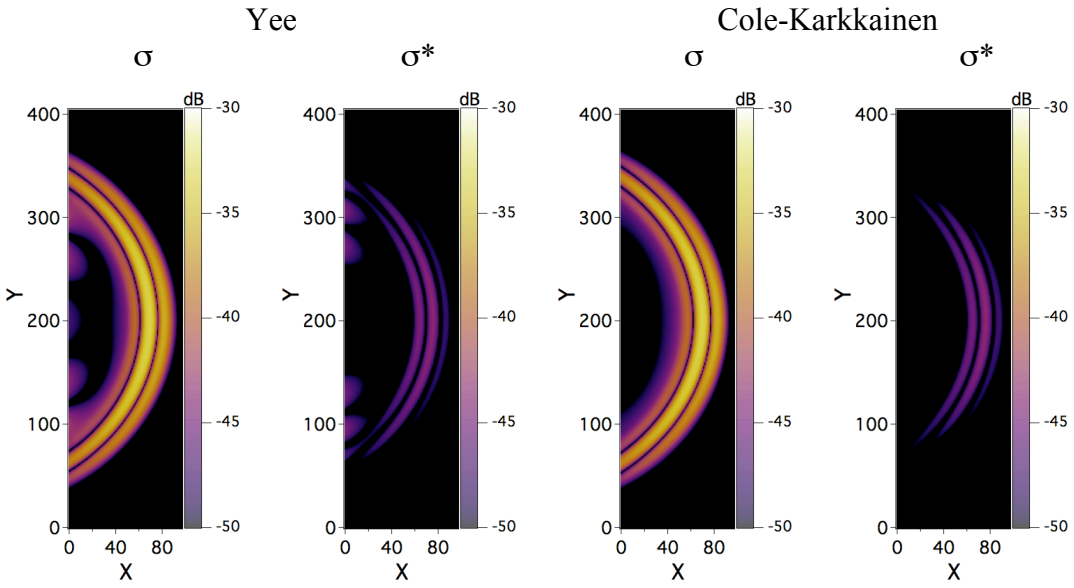

Fig. 2: Reflected signal (in dB) from a PML layer using the Yee or the Cole-Karkkainen solver. Each simulation was run for the time step set at the Courant limit.

For the generic test case that has been considered, the new implementation exhibited a very low residue of reflections from the PML layer, which are qualitatively and quantitatively very similar to the residue obtained with a standard PML implementation. In agreement with results from [13,14], the use of the modified coefficients σ* led to an order of magnitude improvement over the use of the standard coefficients. The 3D tests gave similar absorption efficiency between the Yee and the new solver implementations of the PML.

I-2-b – Unsplit Cummer's PML

We have implemented, in the Warp code's 2-D electromagnetic field solver, an un-split version of the PML algorithm based on Cummer's derivation [15,16], as a complement to the original split implementation of asymmetric PML (APML) [14]. For testing, a Harris pulse was launched at location {X,Y}={200,100} of a $400\delta x \times 400\delta y$

grid, with $L=50\ \delta x$ and $\delta x=\delta y$. The time step δt was set at 99% of the Courant limit and the simulation was stopped when the time, initially set at zero, reached $t_{max}=200\delta x/c$. The electric field at t_{max} is plotted versus x and y in Fig. 3.

The electric field located in the band $0<y<100\delta y$ contains a reflected wave due to the presence of the PML that is revealed by differencing with the free space electric field located in the band $100\delta y <y<200\delta y$, inverted in y. We used a layer 10 cells thick, with numerical conductivity following the progression $\sigma=\sigma_{max}\ (j\delta x/\Delta)$ inside the layer, with $\sigma_{max}=4/\delta x$, $\Delta=5\delta x$, and $n=2$. The result is plotted in Fig. 4 for the newly implemented un-split PML, and is contrasted with the result obtained from the split PML. The un-split PML and split PML give the same amount of reflection, for all angles of incidence.

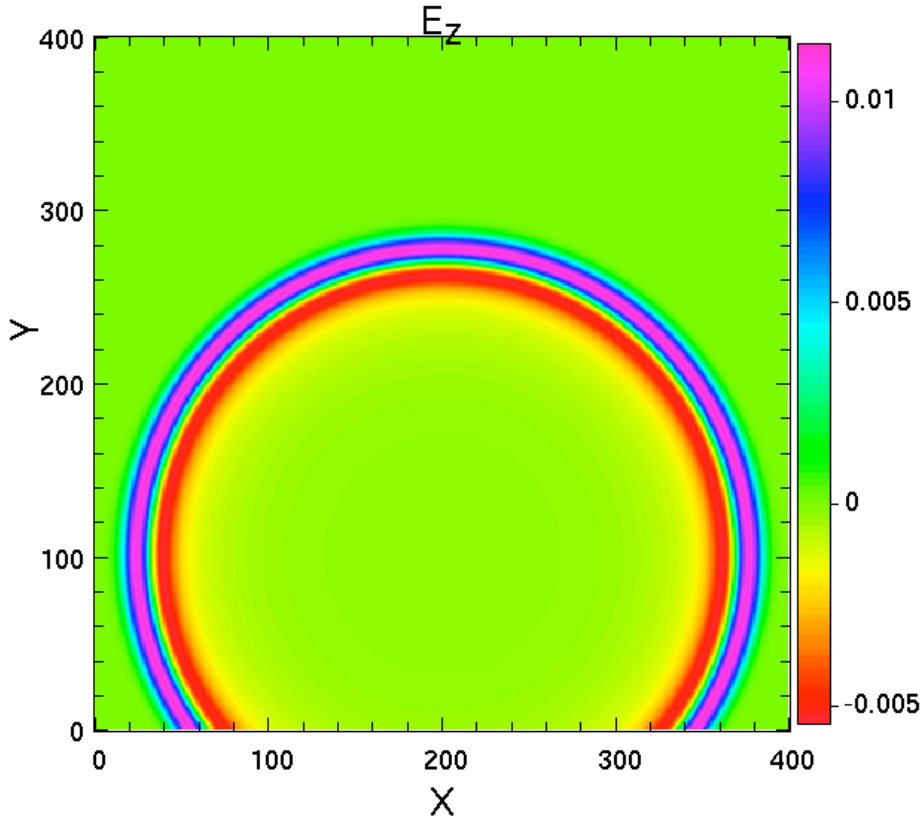

Fig. 3: Electric field from a pulse generated at {x,y}={200,100}

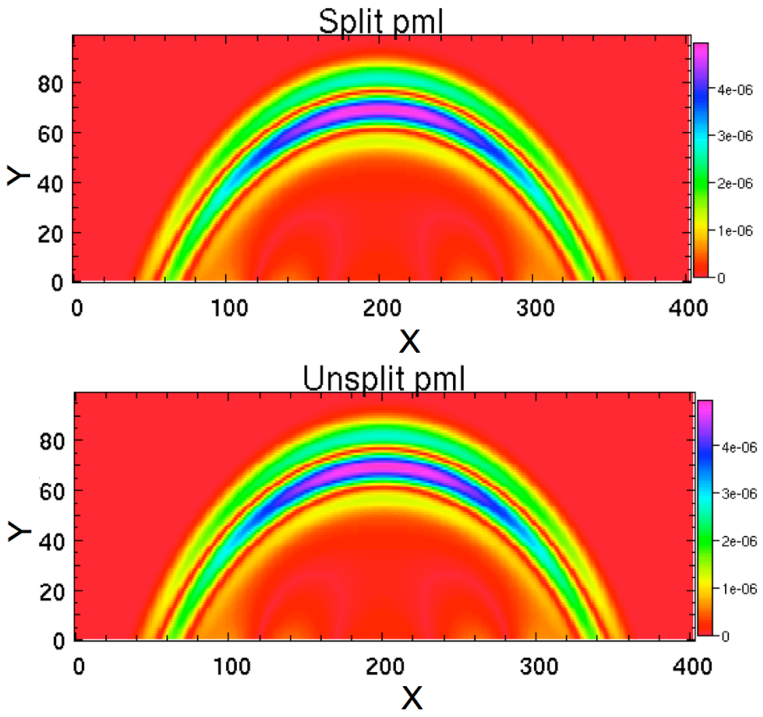

Fig. 4: Electric field reflected by (top) split PML; (bottom) unsplit PML.

I-3 – Mesh refinement

A derived class of the MeshRefinement class (developed for use with the electrostatic/magnetostatic solver class) and the EM3D class was introduced in Warp, providing 3-D electromagnetic mesh refined capability. The implementation follows the scheme based on field substitution ("SG-FS") described in [1]. In order to extend the single level of refinement scheme from [1] to an arbitrary number of refinement levels, an additional auxiliary grid was added, as depicted in Fig. 5.

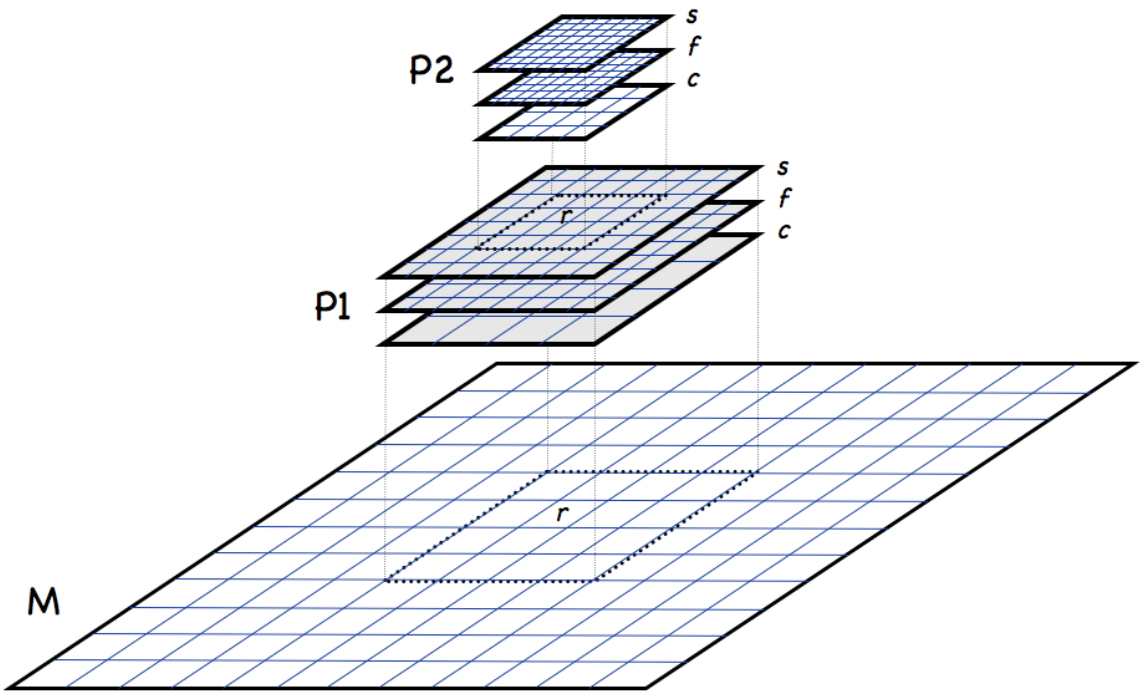

Fig. 5: Schematic of the SG-FS mesh refinement scheme implemented in the Warp 3-D EM-MR solver.

One refinement block comprises three grids: one ("*c*") at the same resolution as the parent grid, and two ("*f*" and "*s*") at the refined resolution. The subset of the parent grid being refined at a given level is denoted by "*r*". The Maxwell equations are solved on *c* and *f*, while the grid *s* is used as an auxiliary grid, storing the field of *f* to which is added the interpolated difference between *r* and *c*, i.e. $s=f+I[r-c]$, where I is the interpolation operator. The purpose of the grid *s* is to store the substituted field that is to be interpolated to the particles, allowing a single lookup per particle for each field component. In practice, PIC calculations are dominated by current deposition and field lookup operations between the particles and the grids, and the extra cost added by the grids *s* should be minimal in most cases.

The current deposition scheme uses the Esirkepov algorithm [11] with the choice at the user-level of linear, quadratic or cubic splines, independently for each particle species and each direction (*x*, *y* or *z*). For linear current deposition, as noted in [1] for particles present in a given refinement patch, the current deposition is performed at the finest level, and subsequently propagated recursively to the parent levels without losing the charge conservation property at any level. For higher-order deposition, however, such a property is not preserved. Although it may be possible to alter the Esirkepov scheme to account for the additional constraint, it goes beyond the scope of the current effort and has not been considered. Thus, in order to use high-order deposition, we employ a modified set of Maxwell equations, as introduced in [17] (and independently by other authors not cited here), which advects residual errors in Gauss Law:

$$\begin{cases} \dfrac{\partial \vec{B}}{\partial t} = -\vec{\nabla} \times \vec{E} \\ \dfrac{\partial \vec{E}}{\partial t} = \vec{\nabla} \times \vec{B} - \vec{J} + \vec{\nabla} F \\ \dfrac{\partial F}{\partial t} = \vec{\nabla} \cdot \vec{E} - \rho \end{cases}$$

Note that solving for the additional term F is optional and controlled at run time.

I-4 – Generalized particle injection

A generalized particle injection capability has been implemented in Warp, to enable injection of particles from any specified conductors. The algorithm is similar to that described in [18]. The number of particles to inject is calculated so as to satisfy Gauss' Law near the surface of the conductors. The integral of the surface-normal E field is calculated over the surface of dual cells, the volume extending from (i-1/2,j-1/2,k-1/2) to (i+1/2,j+1/2,k+1/2). This gives the total charge that should be in the volume. From this is subtracted the charge of the existing particles, that is, the charge deposited at the cell (i,j,k). Enough particles are then injected in the cell to account for the remaining charge. The particles are initially placed directly on the conductor surface and allowed to propagate from there. An illustration of the use of the new injection method is shown in Fig.6.

Currently, the algorithm is implemented using the electrostatic fields, which are calculated using a one-point finite difference stencil, modified to take into account the presence of the conductors. For the electromagnetic version, the E fields used in the integral are coincident with the E fields of the Yee mesh, and so would be a direct drop-in replacement for the electrostatic fields, leaving the rest of the code unchanged.

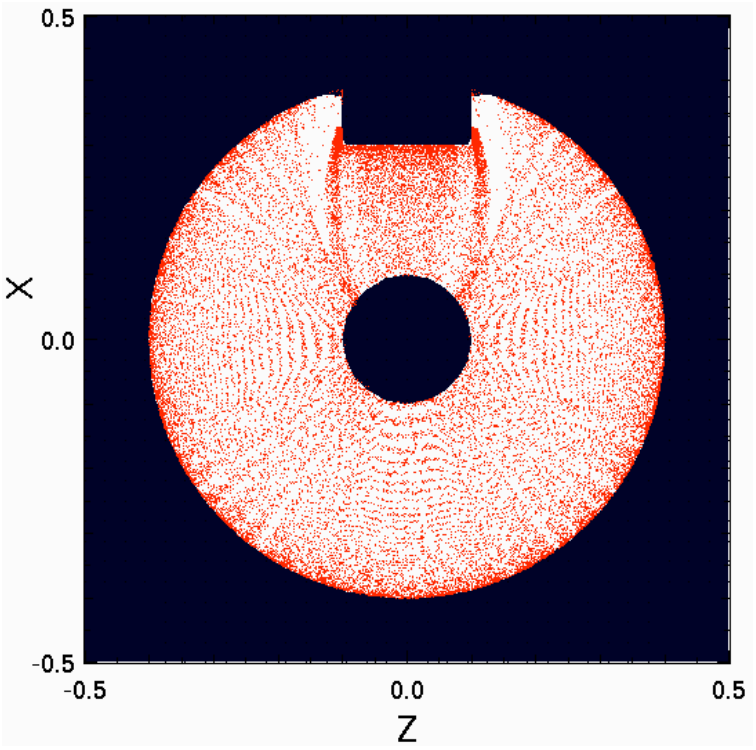

Fig.6: An example demonstrating the Gauss' Law injection. This shows a slice from a simulation of two concentric spheres. The outer sphere has an additional block attached on its inner surface and is at a higher potential relative to the inner sphere. The red dots are the particles created on the outer conductor and are flowing toward the inner. Note the enhanced emission from the edges of the block.

I-5 – Internal conductors

A stair-cased representation of internal perfect conductors was implemented into Warp's 2-D electromagnetic solver, as a first step toward implementing space-charg-limited emission from a shaped surface. The implementation was tested on a hypothetical cavity consisting of a circular core extended by six identical "arms" spaced regularly in azimuth (see Fig. 7). A point-source emitter located at the center of the cavity continuously generated a sinusoidal varying electric field. Interference patterns quickly developed. Without errors due to discretization on a Cartesian mesh, the solution should follow a perfect 6-fold symmetry in azimuth. The interference pattern (snapshot of electric field) shown in Fig. 7 shows that the symmetry is reasonably well respected in the case of a 200x200 grid, but not in the case of a 100x100 grid.

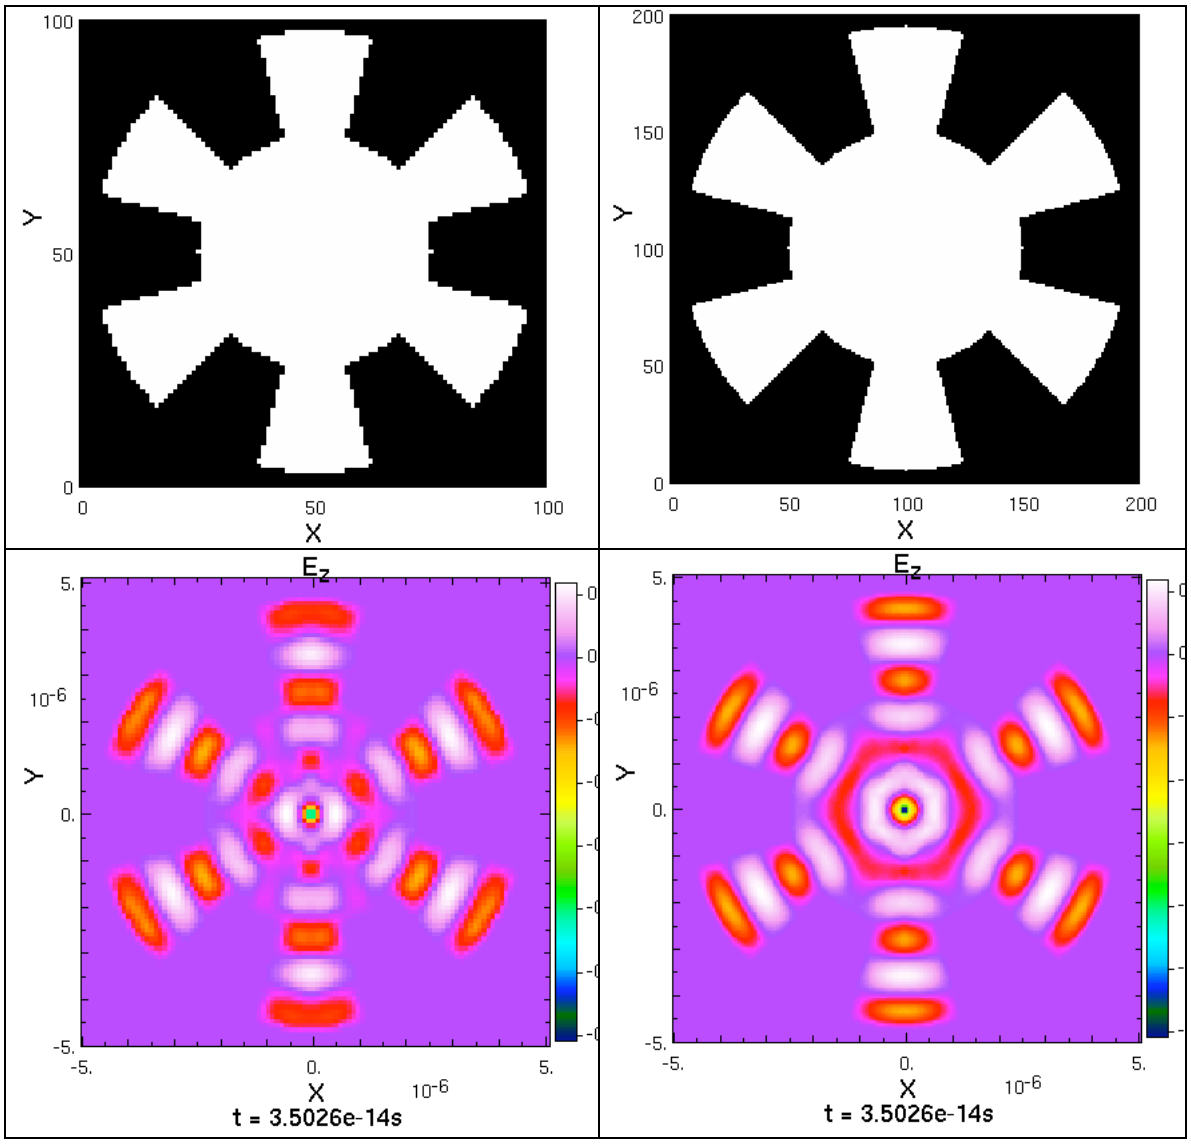

Fig. 7: (top) stair-cased representation of a 6 “arms” cavity for a 100x100 (left) and a 200x200 (right) grid. (bottom) Electric field from a point source emitter located at the center of the cavity.

I-6 – Particle subcycling

We have extended Warp’s sub-cycling [19] capabilities, developed originally for electrostatic PIC simulations, to simulations using the 2-D electromagnetic solver. Each particle species can be advanced using its own time step, which can be different altogether from the electromagnetic field time step. Furthermore, following the implementation for the electrostatic solver, particles from a single species may be advanced with different time steps, dynamically changing time steps based on an instantaneous set of constraints, as described in report 3 of phase I, thereby reproducing a feature of the algorithm for multi-scale Particle-in-Cell plasma simulations proposed in [20].

I-7 – Parallel performance

Warp parallel performance when performing EM-PIC simulations has been characterized on Franklin and other platforms, with scaling to 10's of thousands of processors, as shown in Fig. 8.

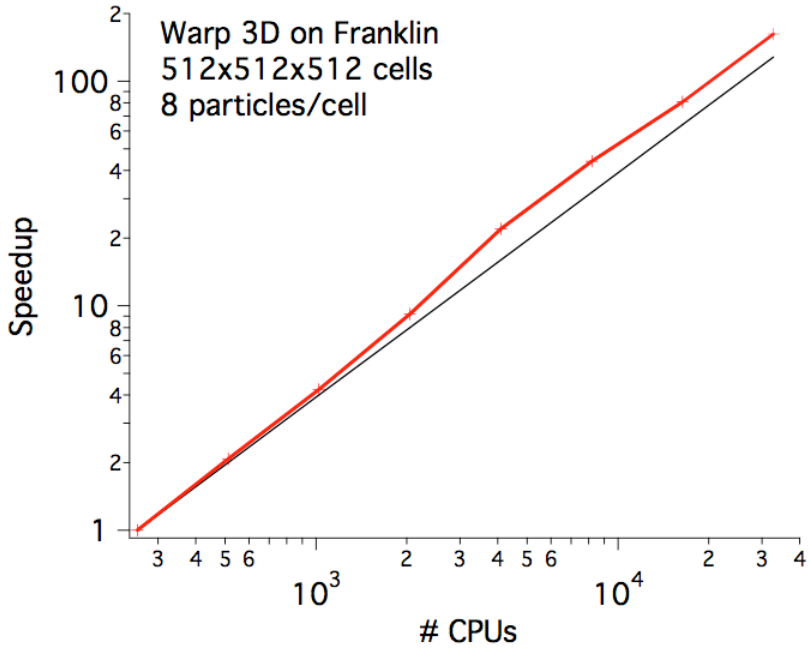

Fig. 8: Parallel speedup of WARP for a strong scaling test of explicit electromagnetic PIC simulations, using a 512x512x512 grid, 8 particles/cell, periodic BC for particles and fields, random load for particles. It was performed on Franklin (NERSC) using from 256 to 32,768 cores. For each run, the timing was averaged over 20 time steps (no large fluctuations were noticed on timings for individual time steps).

2 - Tests and examples of application

2-1 - Filamentation instability

The newly implemented low-dispersion solver was tested on a 2-D planar simulation of filamentation in collaboration with S.M. Lund. The z-axis is perpendicular to the plane of simulation and the boundary conditions are periodic for fields and particles in x and y. The grid had a resolution of 256×256. An exactly charge neutral plasma was initialized with the following prescription: uniform background of density n_i of immobile positively charged particles with infinite mass; two populations of electrons (beam electrons "be" and plasma electrons "pe") of total density $n_{be}+ n_{pe} = n_i$ and $n_{pe}= 10\times n_{be}$ counter-streaming in the direction perpendicular to the plane of calculation, initialized with small random velocities. The initial random velocity load of the electrons creates random fluctuations, seeding a coalescence of the beam electrons into small filaments, which eventually coalesce further until all the beam electrons gather in a large cluster.

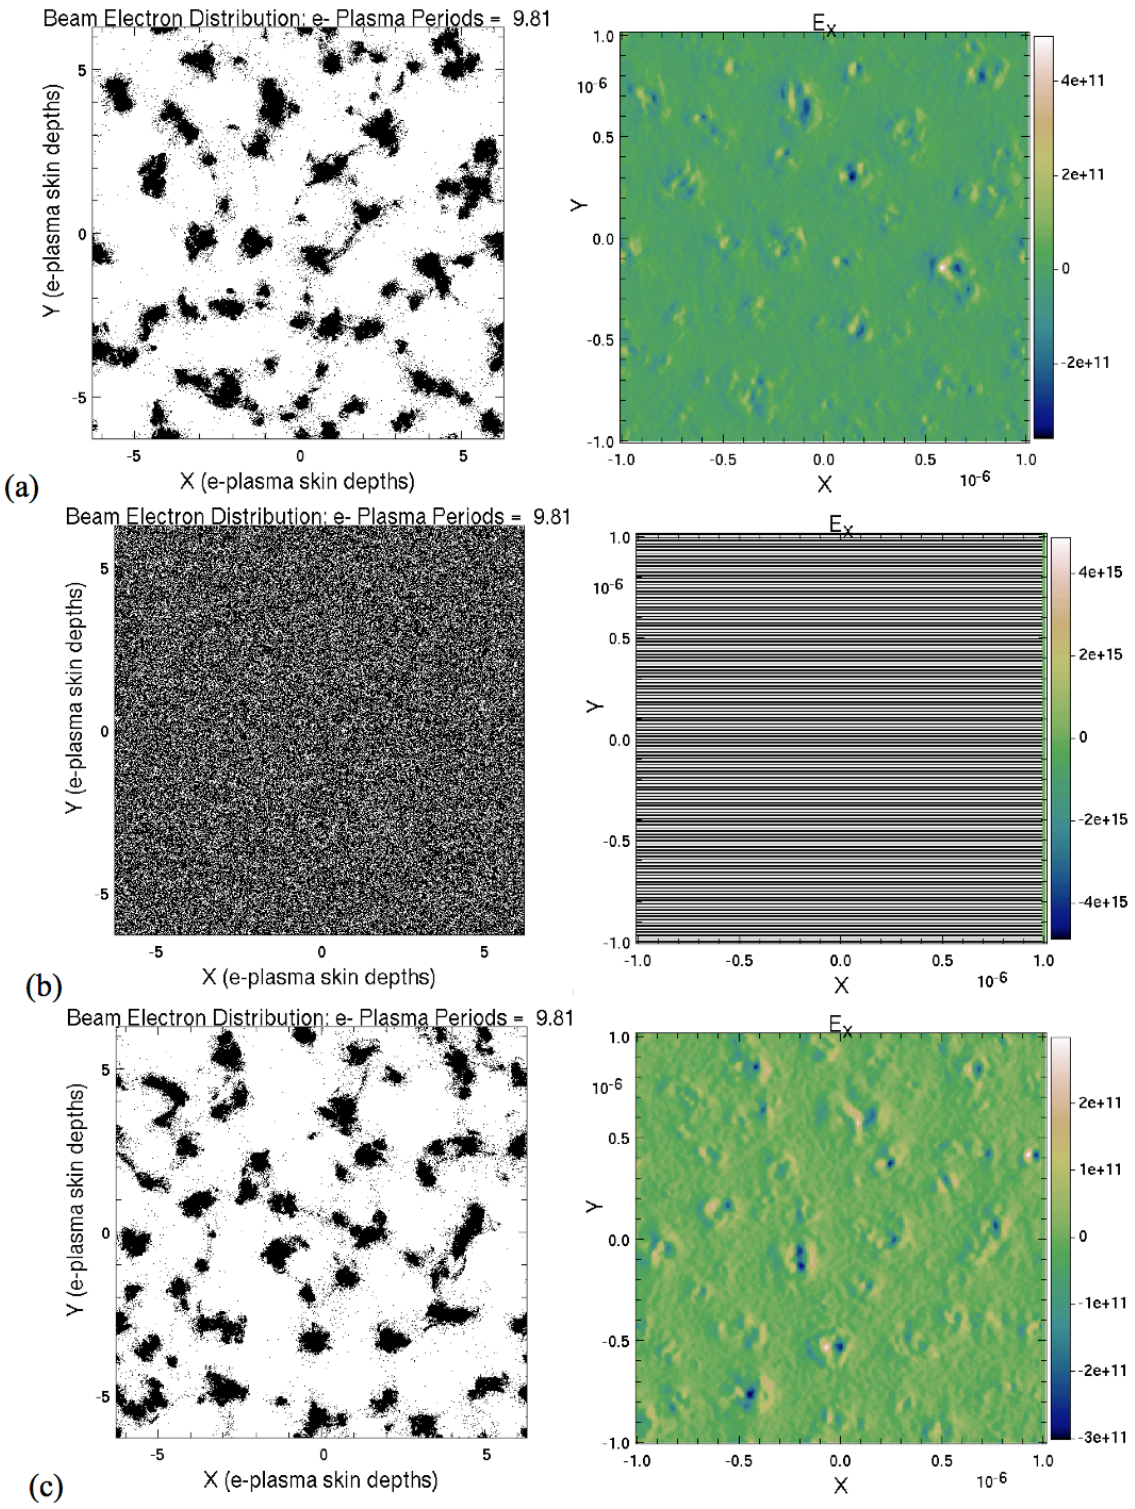

Fig. 9: Snapshots of x-y distribution of beam electrons and the electric field after the simulation has run for almost 10 plasma periods for: a) standard Yee EM solver with $\Delta t=0.98\cdot\Delta x/(\sqrt{2}\cdot c)$, b) low-dispersion EM solver with $\Delta t=\Delta x/c$, and c) low-dispersion EM solver with $\Delta t=0.99\cdot\Delta x/c$.

The x-y distribution of beam electrons and the electric field map are given in Fig. 9 after the simulation has run for almost 10 plasma periods for three cases: (a) standard Yee EM solver with $\Delta t=0.98\cdot\Delta x/(\sqrt{2}\cdot c)$; (b) low-dispersion EM solver with $\Delta t=\Delta x/c$; and (c) low-dispersion EM solver with $\Delta t=0.99\cdot\Delta x/c$. We observe that runs (a) and (c) gave essentially the same level of filamentation at the same physical time while run (b), the one where the time-step used was exactly the magic time-step $\Delta t=\Delta x/c$, exhibits no filamentation, but exhibits signs of odd-even instability in the y direction for the electric field component E_x. The energy plot given in Fig. 10 confirms what seems to be an exponential growth of the field energy for run (b), while the energy histories look very similar for (a) and (c).

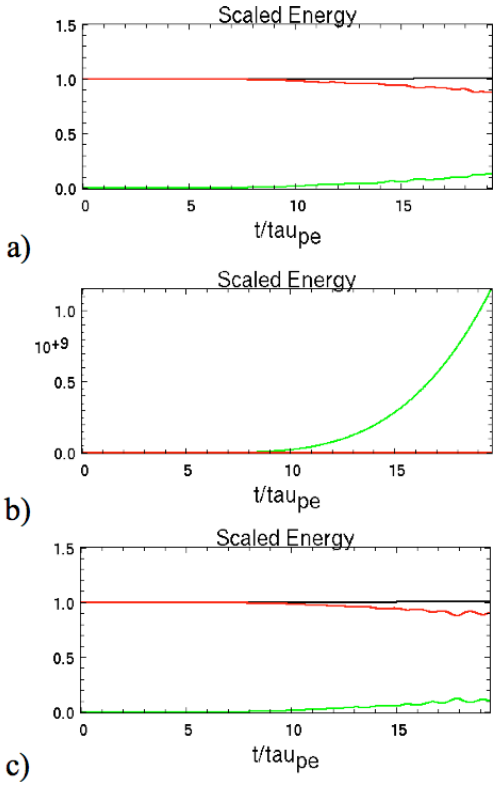

Fig. 10: Energy history for particles (red), fields (green) and total (black) for : a) standard Yee EM solver with $\Delta t=0.98\cdot\Delta x/(\sqrt{2}\cdot c)$, b) low-dispersion EM solver with $\Delta t=\Delta x/c$, and c) low-dispersion EM solver with $\Delta t=0.99\cdot\Delta x/c$.

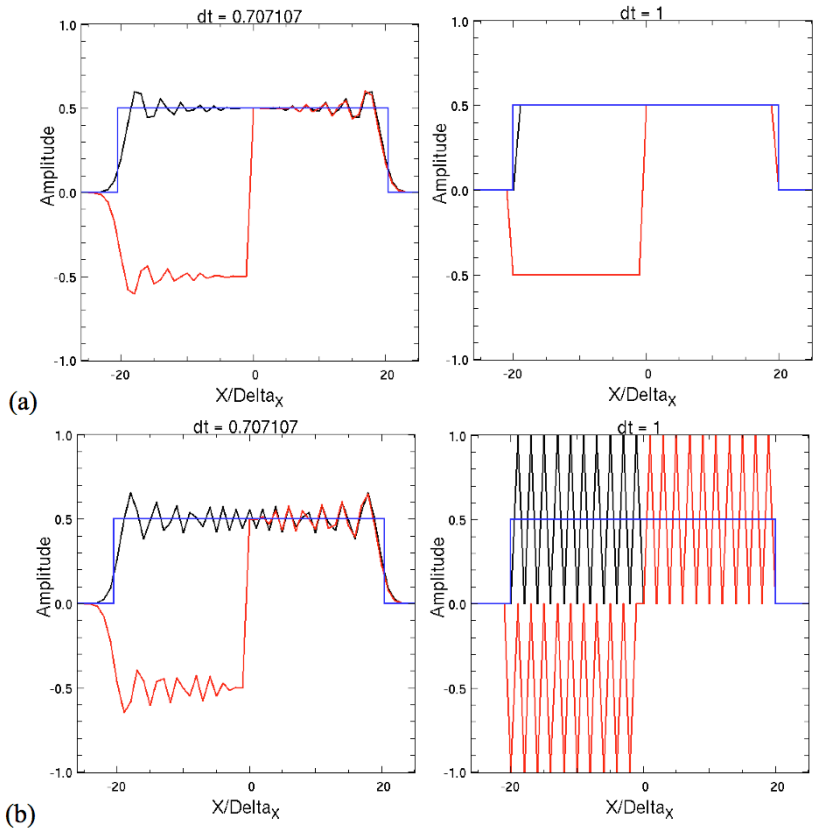

Fig. 11: Snapshot of electric (black), magnetic (red) and analytical electric (blue) fields given from solving the 1-D leapfrog wave equation (a) without source term (the electric field is forced to a constant value at x=0); (b) with a source term, for $\Delta t=\Delta x/(\sqrt{2}\cdot c)$ (first column) and $\Delta t=\Delta x/c$ (second column).

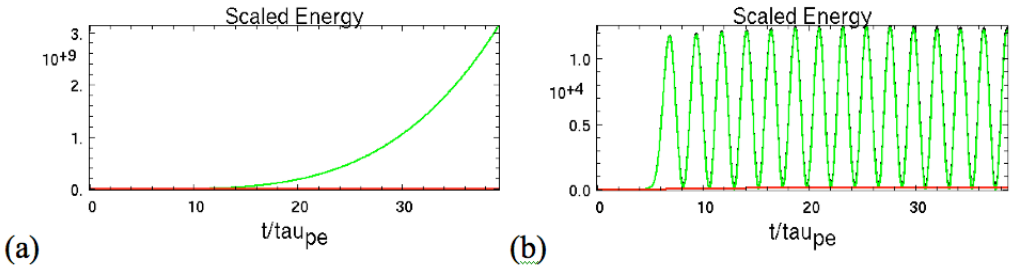

Fig. 12: Energy history for particles (red), fields (green) and total (black) for a run with the low-dispersion EM solver with $\Delta t=\Delta x/c$, for a grid of (a) 64×64, and (b) 65×65.

The odd-even pattern observed in Fig.9 is reminiscent of the odd-even pattern (i.e. signal at the Nyquist limit) obtained when solving the Leapfrog discrete wave equation in 1-D at the Courant limit ($\Delta t=\Delta x/c$) with a constant source term, as shown in Fig. 11. Assuming that the odd-even oscillation of the 1-D case is the source of the instability observed in the 2-D run, one would expect that changing the number of grid cells from

even to odd would change the way that interferences arising from the use of periodic boundary conditions will add constructively or destructively over the course of a run. We thus ran a scaled down version of the previous 2-D runs for two griddings: (a) 64×64 and (b) 65×65, and report the respective energy histories in Fig. 12. We effectively observe a very different pattern between the even (a) and the odd (b) cases, with an exponential growth in (a), suggesting continuous constructive addition of the odd-even oscillations, while (b) gives periodic oscillations of the field energy, suggesting that the individual oscillations go periodically in and out of phase.

The instability was successfully mitigated by smoothing the deposited current with a bilinear digital filter, which totally removes signals at the Nyquist frequency. More details are given in [12].

Fig. 13 shows snapshots from Warp simulations of the filamentation instability in 3D.

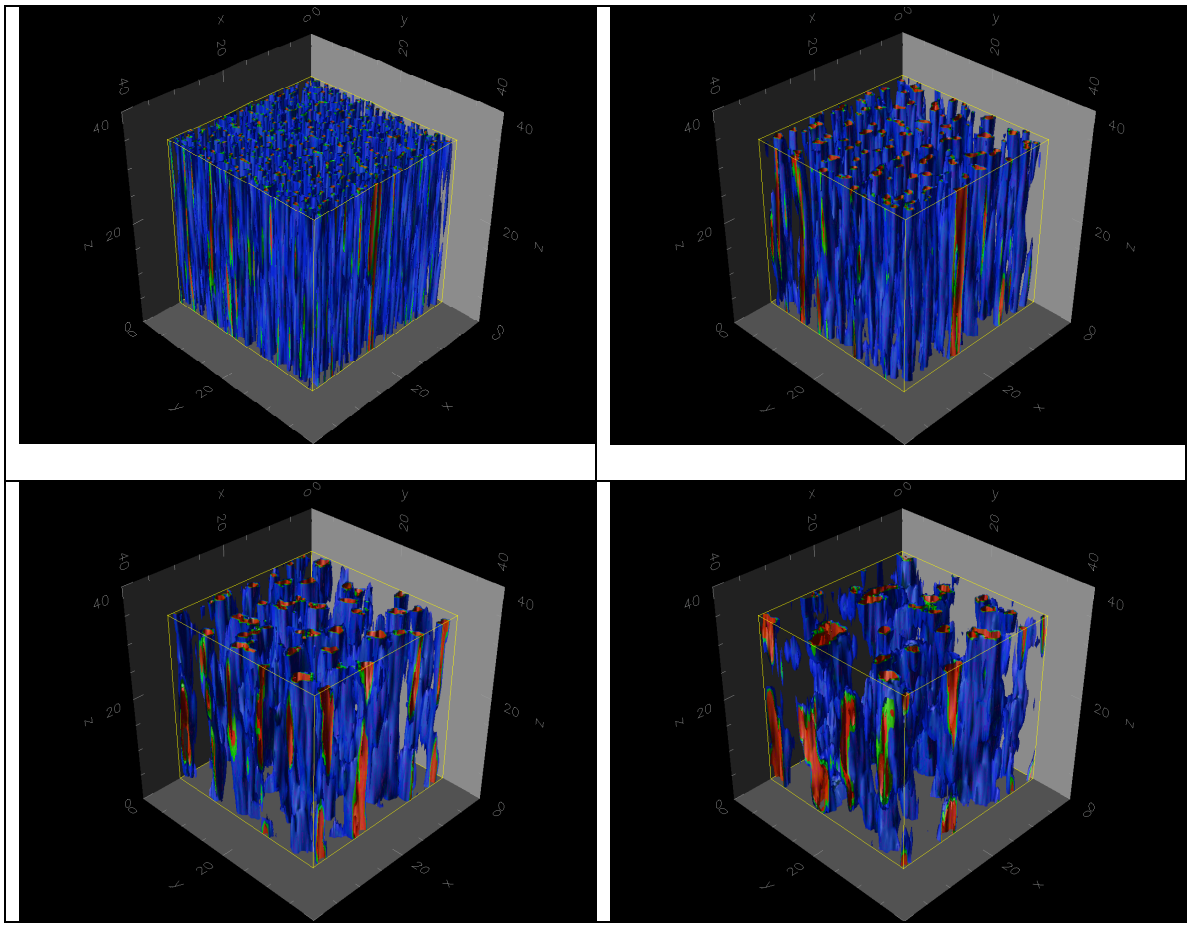

Fig. 13: snapshots of density contours (blue=low density; red=high density) for a Warp 3D run using a 128x128x8 grid on a filamentations instability.

2-2 – Fast ignition

The configuration reported in [1] (see also report 1 of phase I) is considered, of a laser impinging a target above the critical density for fast ignition study.

The parameters are the same as those given in [1], except for the high-density core which was omitted. Two runs were performed for testing the particle subcycling capability. One used a single time step δt for the field solver and both species, at 0.6 times the Courant condition for the fields. The second run used the same time step δt for the field solver, and four time steps for the particles: δt, 2×δt, 4×δt and 8×δt. The criteria for promoting/demoting particles to groups pushed with respectively larger/smaller time-steps was based on thresholds for particle motion across a cell in any direction. These thresholds were set empirically to 1/6 (for promotion) and 1/12 (for demotion) of the cell sizes. Fig.14-18 show snapshots of various projections of the ion and electron distributions taken at T=200ps. The two runs gave similar results. However, the run with adaptive time-steps was faster by about a factor of two (see Fig. 19). The history of the number of macro-particles for each of the four time-step groups is shown in Fig. 20.

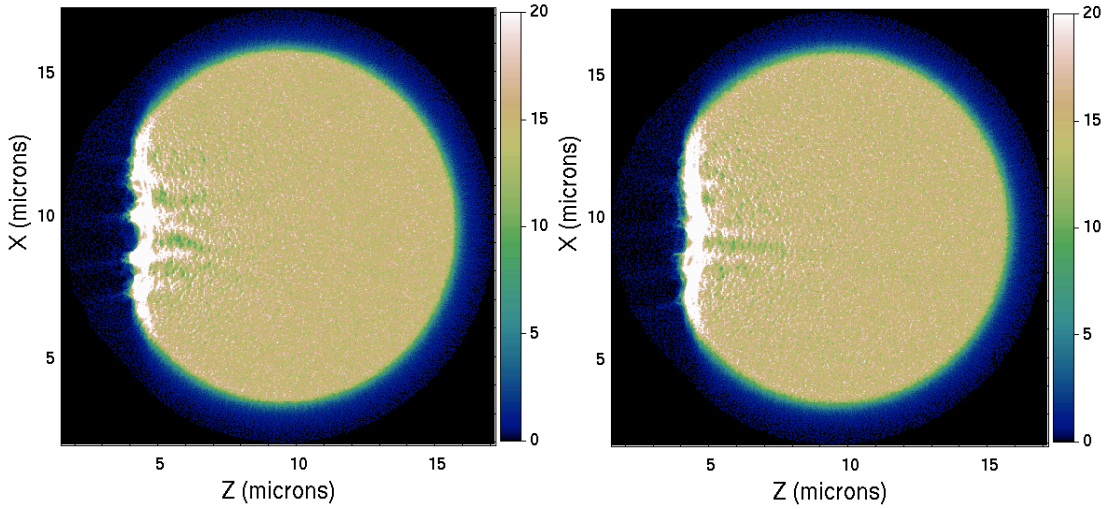

Fig 14: Snapshot of ion distribution in ZX colored according to density, for run with single small time steps (left) and run with adaptive time stepping (right).

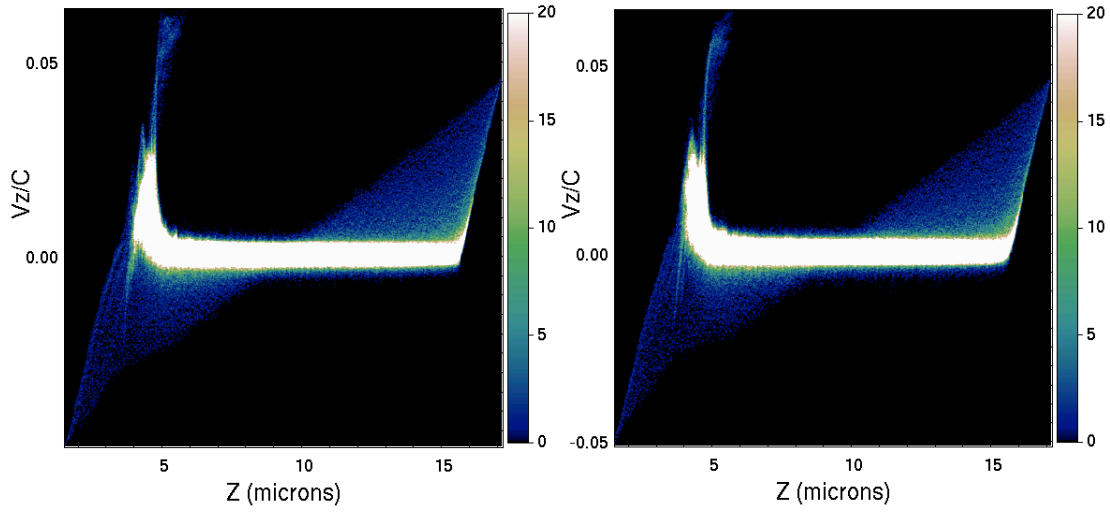

Fig 15: Snapshot of ion distribution in ZVz colored according to density, for run with single small time steps (left) and run with adaptive time stepping (right).

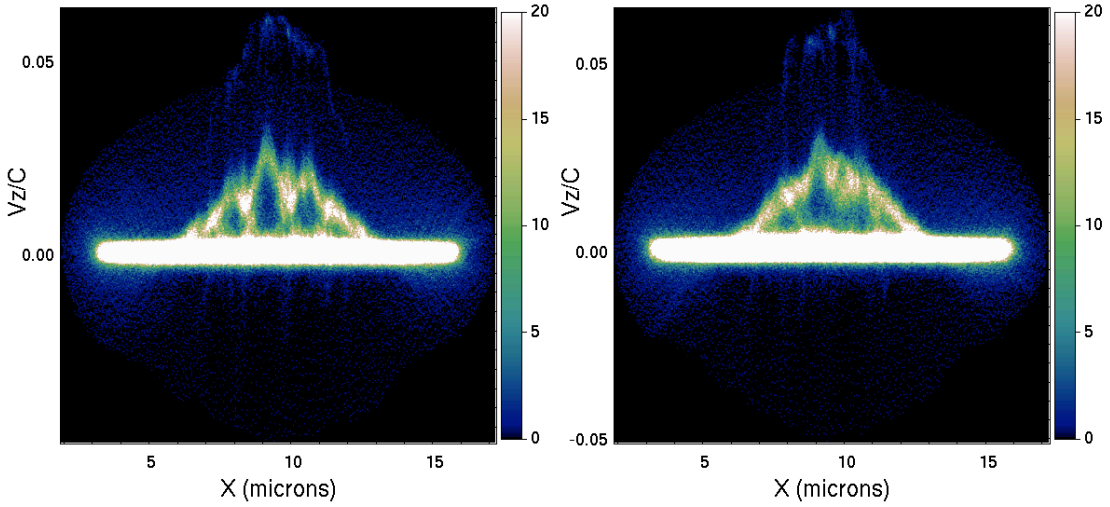

Fig 16: Snapshot of ion distribution in XVz colored according to density, for run with single small time steps (left) and run with adaptive time stepping (right).

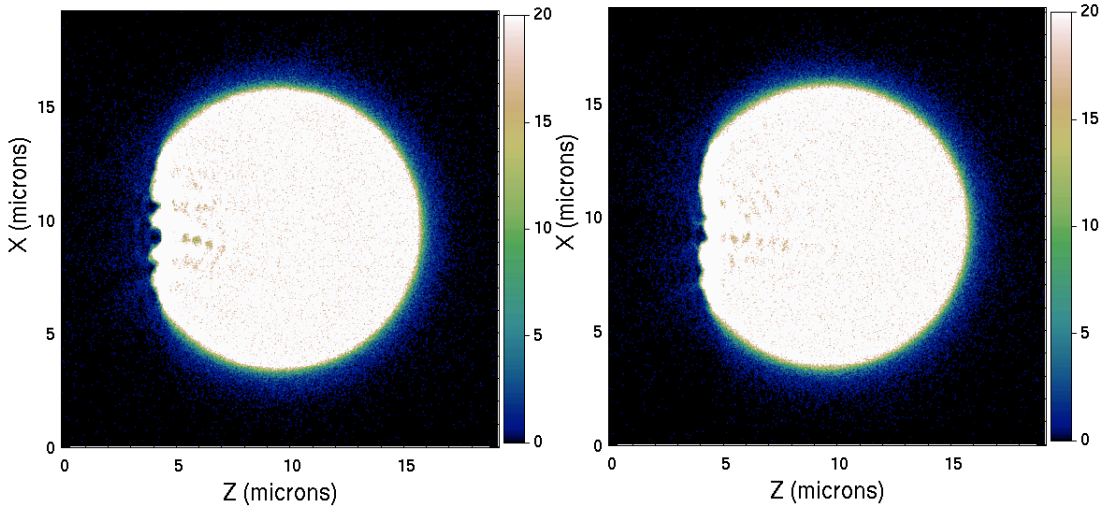

Fig 17: Snapshot of electron distribution in ZX colored according to density, for run with single small time steps (left) and run with adaptive time stepping (right).

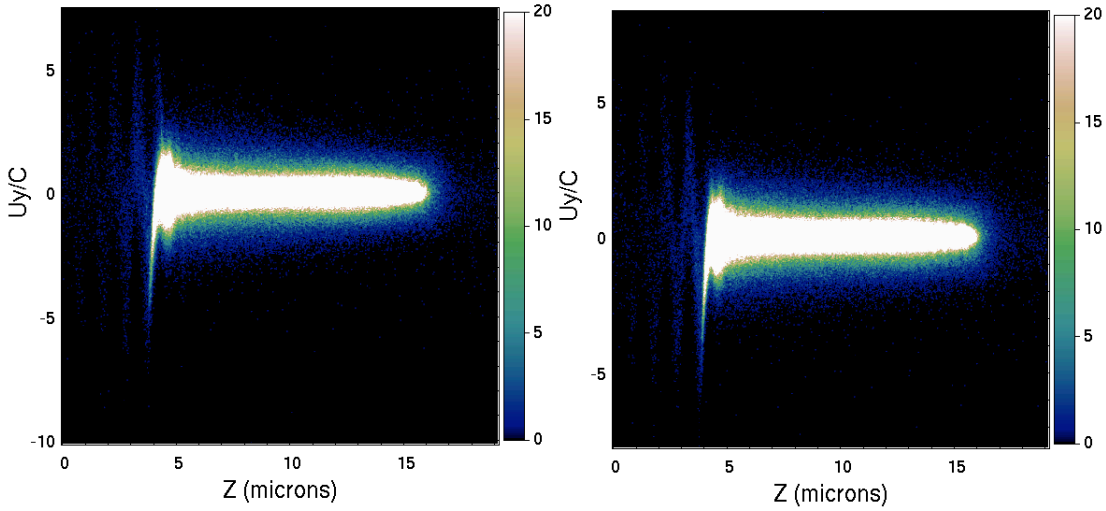

Fig 18: Snapshot of electron distribution in ZUy colored according to density, for run with single small time steps (left) and run with adaptive time stepping (right).

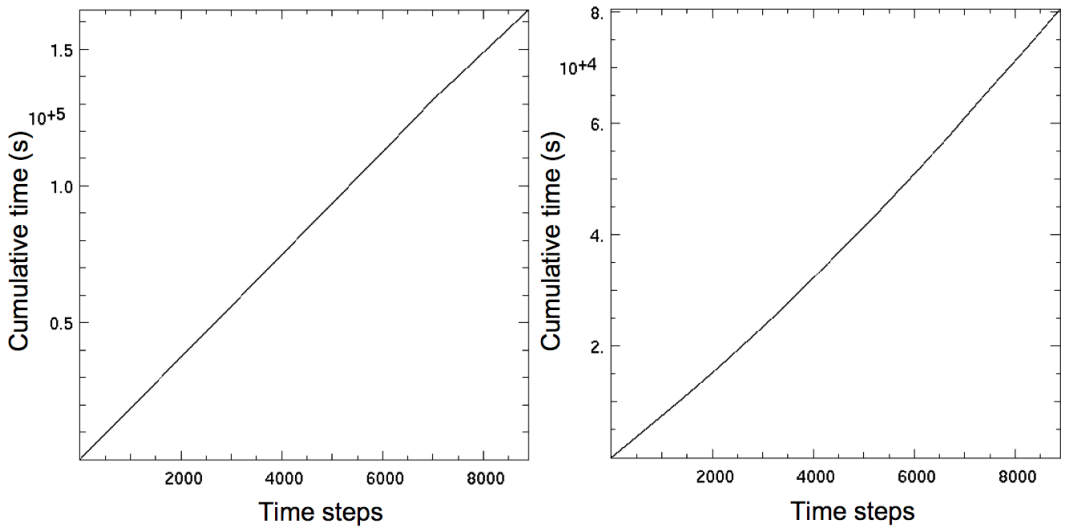

Fig. 19: Cumulative time versus number of time steps, for run with single small time steps (left) and run with adaptive time stepping (right).

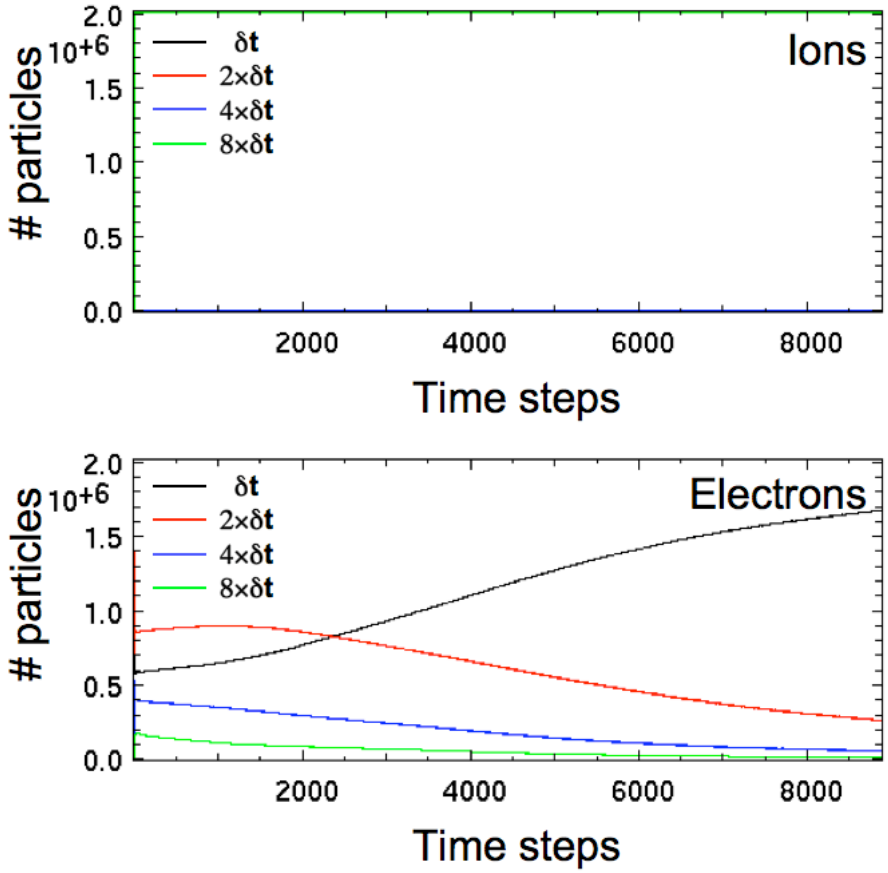

Fig. 20: History of the number of macro-particles for each time-step group, for ions (top) and electrons (bottom).

In [1], it was observed that the mesh refinement technique that we introduced based on grid substitution, may suffer from spurious transmission of waves across patches (in cases where a high-density plasma fills the patch and normally absorbs the waves in consideration). This spurious effect was attributed to inexact cancellation of waves during the substitution procedure, resulting from the numerical dispersion errors of the electromagnetic solver.

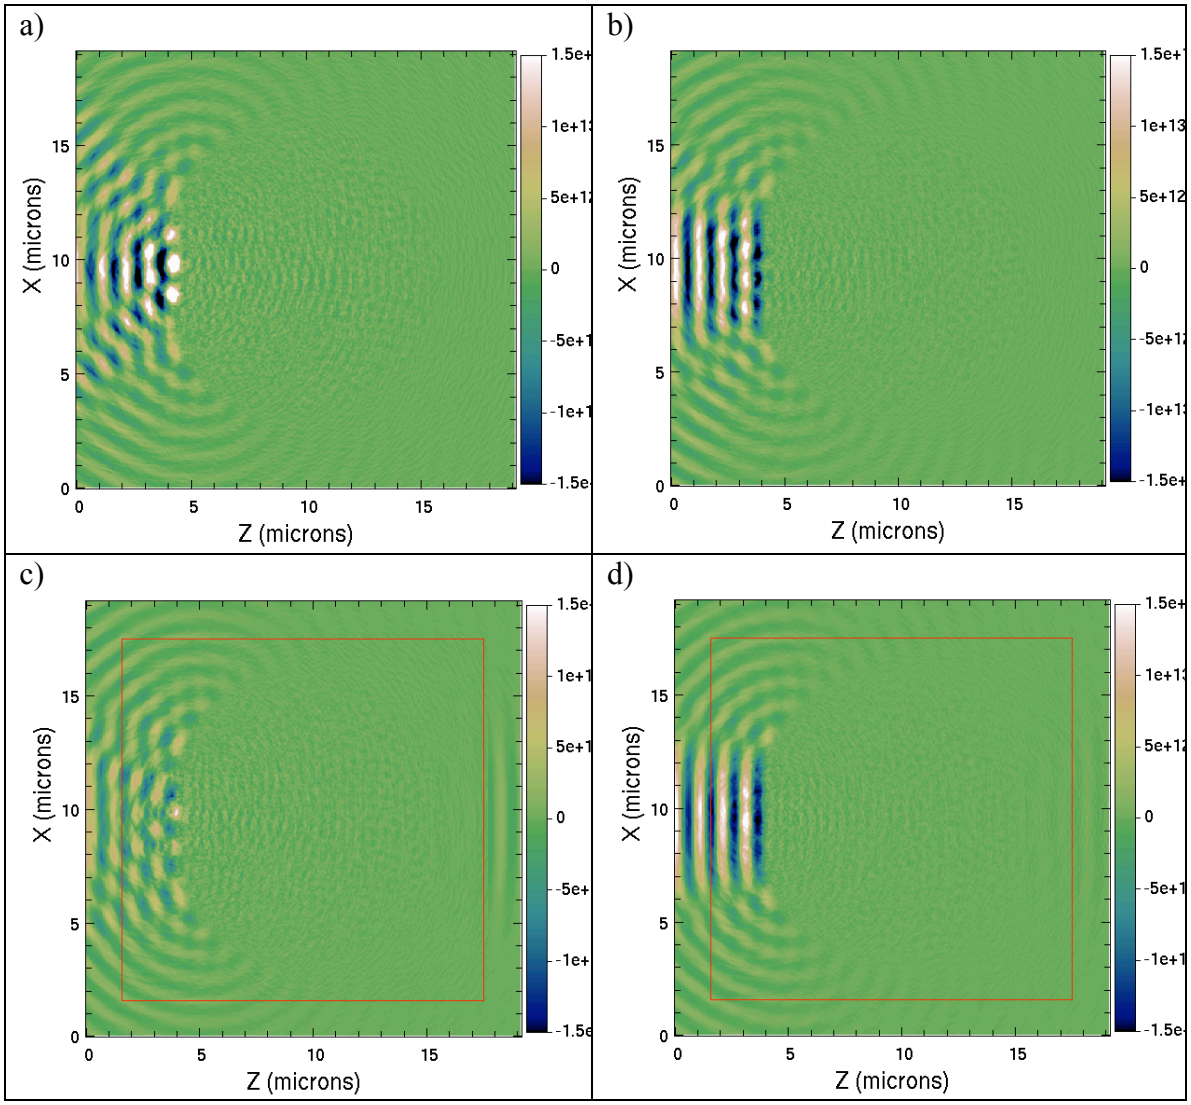

Fig. 21: Colored plot of magnitude of transverse electric fields from four Warp runs: a) 600*600 grid/no mesh refinement/Yee stencil; b) 600*600 grid/no mesh refinement/ Karkkainen stencil; c) 300*300 grid/one mesh refinement patch/Yee stencil; d) 300*300 grid/one mesh refinement patch/ Karkkainen stencil. The red line in plots c) and d) delimits the boundary of the refinement patch.

To test this hypothesis, four runs were performed using Warp 3-D solver in 2-D1/2 mode: a) 600*600 grid/no mesh refinement/Yee stencil; b) 600*600 grid/no mesh

refinement/Cole-Karkkainen stencil; c) 300*300 grid/one mesh refinement patch/Yee stencil; d) 300*300 grid/one mesh refinement patch/Cole-Karkkainen stencil. For the cases with refinement patch, the refinement factor was 2 in each direction, and the time step of the main grid was twice the one used to push particles and fields in the refinement patch (interpolation in time is used to get fields from the main grid at intermediate time steps during the substitution procedure). The electric field transverse to the plane of calculation is plotted in 2-D (Fig. 21) and in 1-D (cut at x=9.55μm; Fig. 22). The spurious wave described in [1] is observed here also, but its amplitude is much reduced with the use of the CK solver, supporting the hypothesis that it is a consequence of the numerical dispersion errors of the electromagnetic solver.

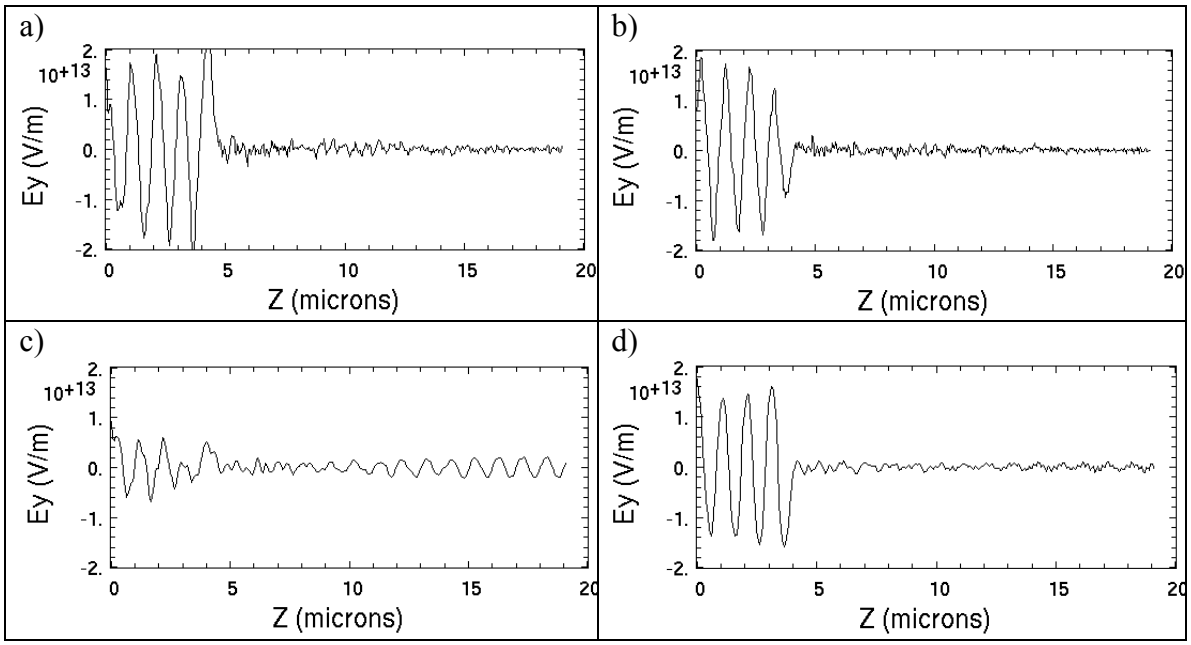

Fig. 22: a) Line plot of magnitude of transverse electric fields (cut at x=9.55μm) from four Warp runs: 600*600 grid/no mesh refinement/Yee stencil; b) 600*600 grid/no mesh refinement/ Karkkainen stencil; c) 300*300 grid/one mesh refinement patch/Yee stencil; d) 300*300 grid/one mesh refinement patch/Cole-Karkkainen stencil.

2-3 - Modeling of ion beam induced plasma wake

Warp simulations of wave excitations by a beam propagating through plasma, as described in [21], were conducted. In these simulations, a hard-edged, elliptical, "frozen" (rigid) beam propagates at constant velocity $v_z = 0.5c$ through an initially cold neutral plasma of initial density n_0. The beam has a flat-top density profile of $n_b = n_0/2$, and an elliptical shape of length $l = 15\ c/\omega_p$ and diameter $d = l/10$, where ω_p is the electron plasma frequency. It is shown in [21] that waves with a wavenumber of approximately $2\omega_p/v_z$ are generated in the plasma by the beam's electrostatic field, and have larger amplitude inside the beam, due to their interaction with the beam sharp edges.

Resolving the beam edge and the small structures developing in the wake inside the beam imposes small cell sizes. The resolution that is needed for macroscopic convergence was explored in 2-1/2D in a series of four runs where the number of grid cells was varied from 64×160 to 512×1280 by incremental factors of 2. Third order spline interpolation was used for the beam and plasma macroparticle current deposition and force gathering. The results are shown in Fig. 23. The details of the plasma wake are very similar between the two highest resolution cases, indicating that macroscopic convergence was reached. The results from the highest resolution run serve as the reference for subsequent calculations with mesh refinement.

A series of three runs were conducted, where the main grid had 128×320 cells, and was complemented by two refinement patches (with successive refinement of 2 in each direction), such that the resolution in the central patch matched the resolution of the case of reference. Results are plotted in Fig. 24. The runs from Fig. 24-a) and 24-b) differ only by the use of mesh refinement in the latter. We observe that the mesh refinement algorithm did not introduce any detectable spurious effect. Although there is some difference in the electron density outside the area of the highest density patch, the result inside the patch is essentially the same as the one from the reference case. The number and weight of the injected plasma macroparticles was varied, such that the number of macroparticles per cell in each grid at injection was constant (Fig. 24-c). Again, the observable differences in the plasma density outside the central refinement patch did not lead to visible alteration of the result in the central refined area. Finally, as a proof that the plasmas and fields present in the low resolutions regions provide essential boundary conditions to the physics in the high resolution central core, a run was conducted where no plasma was injected outside the central area (Fig. 24-d).

Lastly, three-dimensional simulations with mesh refinement of the same physical setup were conducted. The grid setup and 3-D isosurfaces of the plasma electron density as the beam enters the plasma are shown in Fig. 25. As expected, structures similar to the ones observed in 2-D are present within the beam envelope. Those calculations were possible at a reasonable computational cost only with the use of mesh refinement, with the computing resources that were of available to the user at the time of this calculation. The speedup achieved by the use of mesh refinement is estimated to be of approximately one order of magnitude.

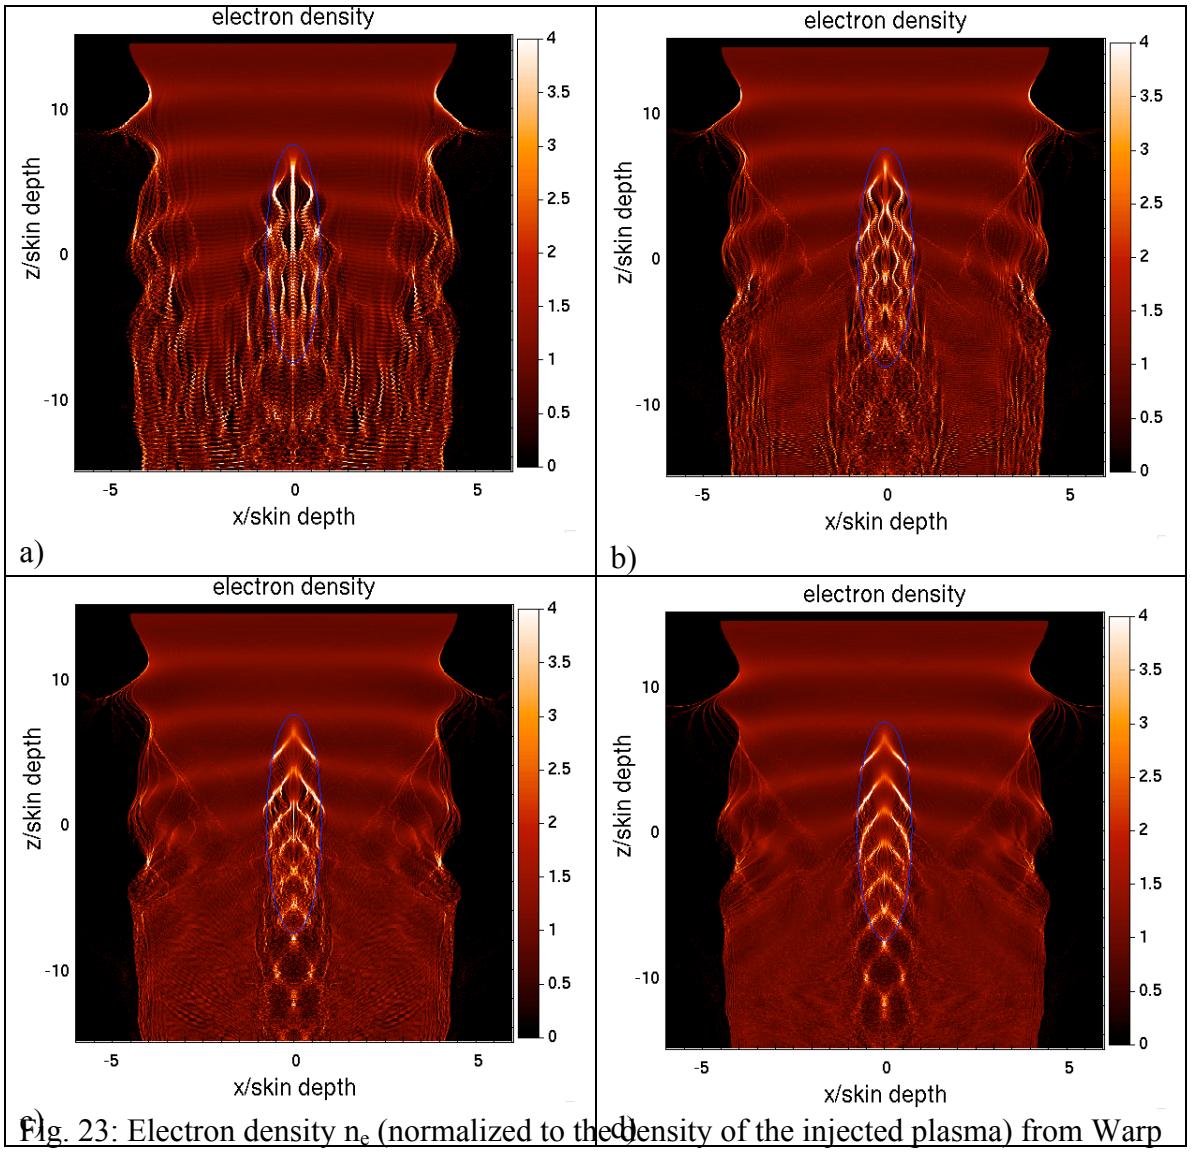

Fig. 23: Electron density n_e (normalized to the density of the injected plasma) from Warp PIC simulations using the 3-D electromagnetic solver in 2-D "slab" (x,z) geometry, of a rigid beam (thin outline) propagating through a neutral plasma, for grid sizes $N_x \times N_z$ of **a)** 64×160, **b)** 128×320, **c)** 256×640 and **d)** 512×1280.

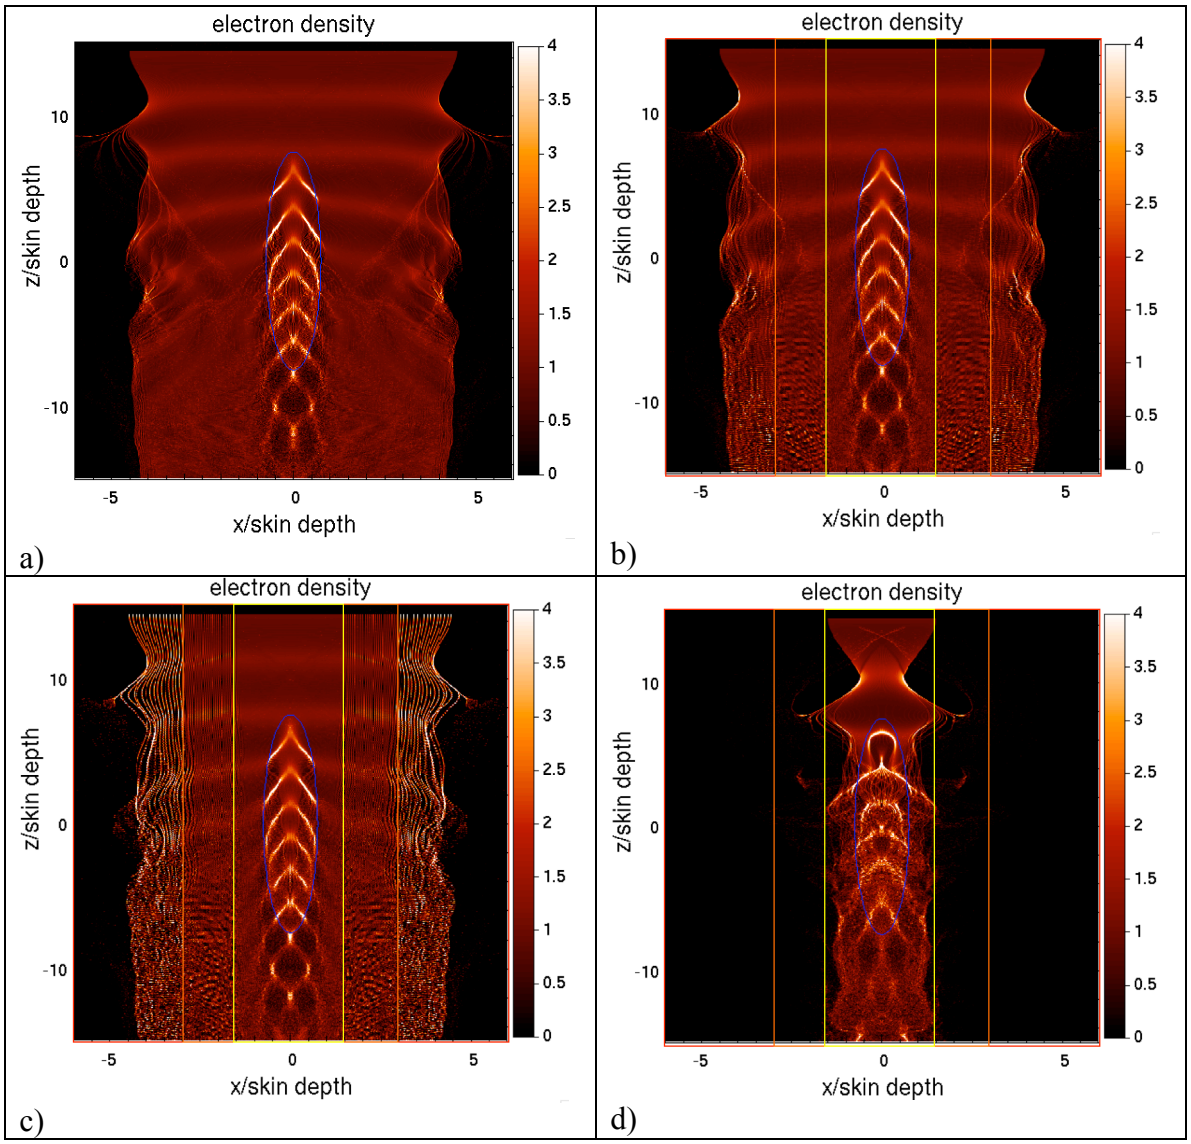

Fig. 24: Electron density n_e (normalized to the density of the injected plasma) from Warp PIC simulations using the 3-D electromagnetic solver in 2-D "slab" (x,z) geometry, of a rigid beam (thin outline) propagating through a neutral plasma, for grid sizes $N_x \times N_z$ of **a)** 512×1280 (same as Fig.1-d), **b-c-d)** 128×320 (main grid, red box) + 128×640 (patch 1, orange box) + 128×1280 (patch 2, yellow box), such that the resolution of patch 2 matched the resolution of the grid used for a). The number and weight of injected plasma macroparticles were: **a-b)** uniform, **c)** adjusted to keep the number of macroparticles per cell constant in each grid at injection, **d)** same as a-c in patch 2 and null outside.

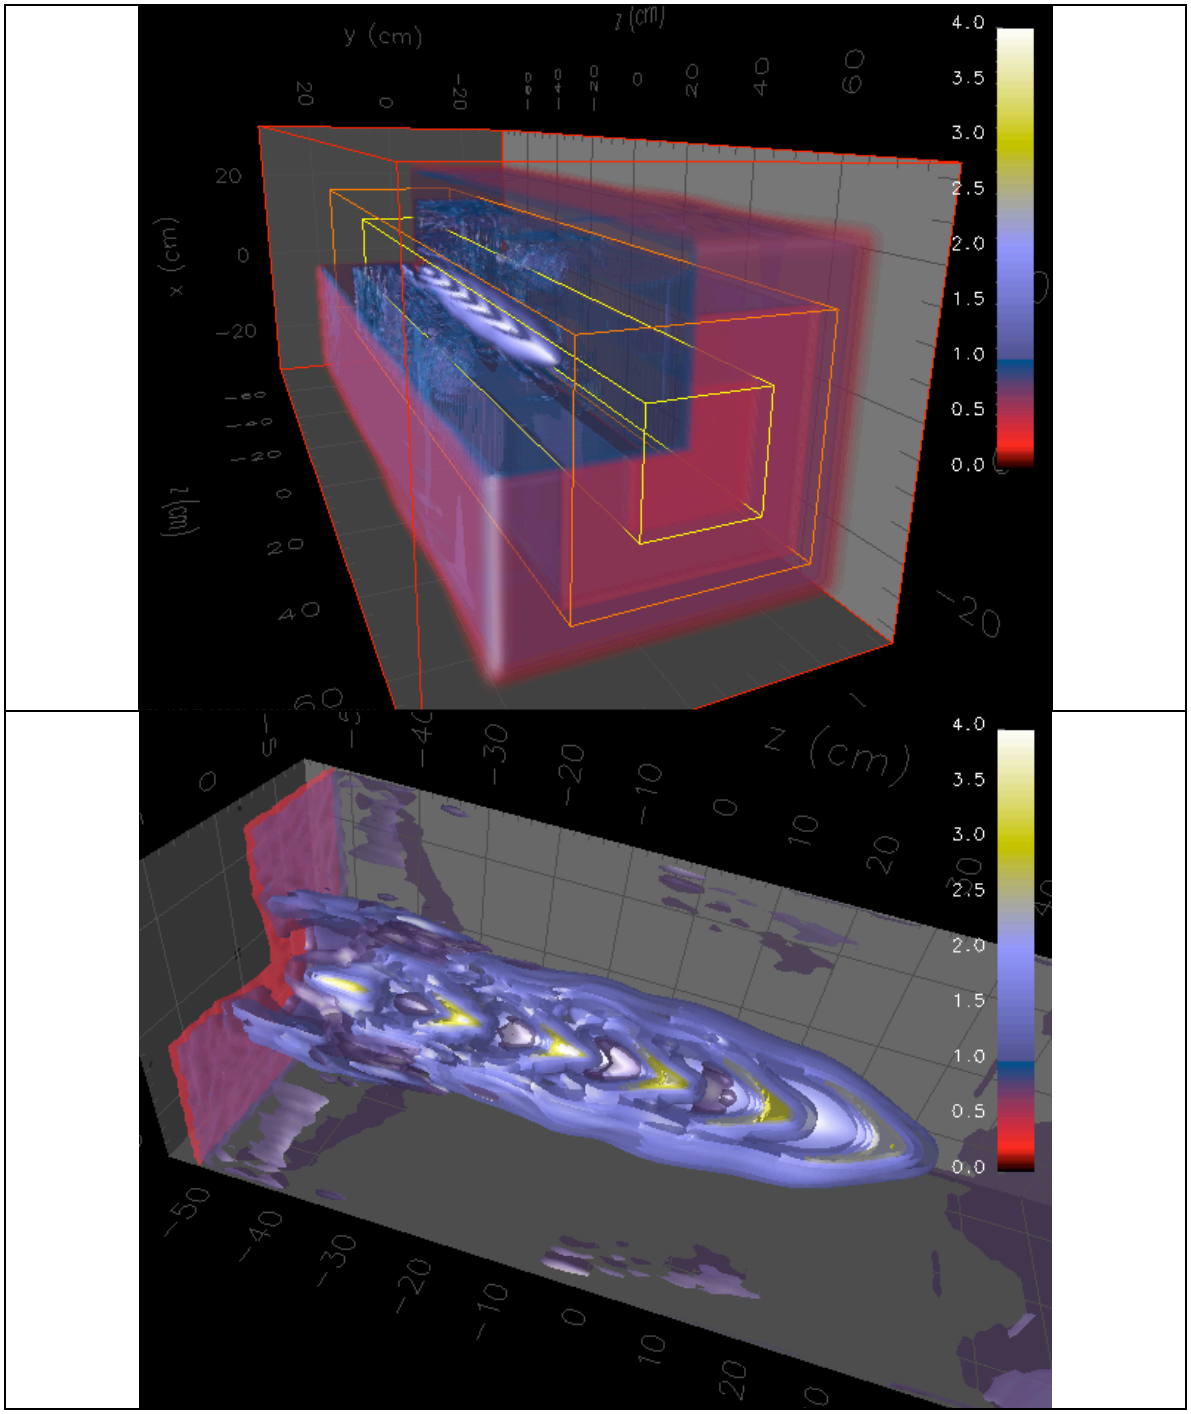

Fig. 25: (top) Electron density n_e (normalized to the density of the injected plasma) from 3-D Warp PIC simulations of a rigid beam propagating through a neutral plasma, for grid sizes $N_x \times N_y \times N_z$ of 64×64×160 (main grid, red box) + 64×64×320 (patch 1, orange box) + 64×64×640 (patch 2, yellow box); (bottom) same as top zoomed at location surrounding the beam.

2-4 - Laser plasma acceleration

Warp's Yee and Cole-Karkkainen solvers were applied in 2-1/2D and 3D to the modeling of a scaled 10GeV stage of a laser wakefield accelerator [22-24], in collaboration with E. Cormier-Michel and C. Geddes from the LOASIS group at LBNL.

Laser driven plasma waves offer orders of magnitude increases in accelerating gradient over standard accelerating structures (which are limited by electrical breakdown), thus holding the promise of much shorter particle accelerators. Yet, computer modeling of the wake formation and beam acceleration requires fully kinetic methods and large computational resources due to the wide range of space and time scales involved [22]. As discussed in [25], the range of scales can be greatly reduced by performing the calculation in a Lorentz boosted frame if one adopts the common assumption that the backward emitted radiation can be neglected.

Figure 26 shows surface renderings of the longitudinal electric field as the beam is in its early stage of acceleration by the plasma wake from the calculation in the laboratory frame and in the frame γ=10. The two snapshots offer strikingly different views of the same physical processes: in the laboratory frame, the wake is fully formed before the beam undergoes any significant acceleration and the imprint of the laser is clearly visible, while in the boosted frame calculation, the beam is accelerated as the plasma wake develops, and the laser imprint is not visible on the snapshot.

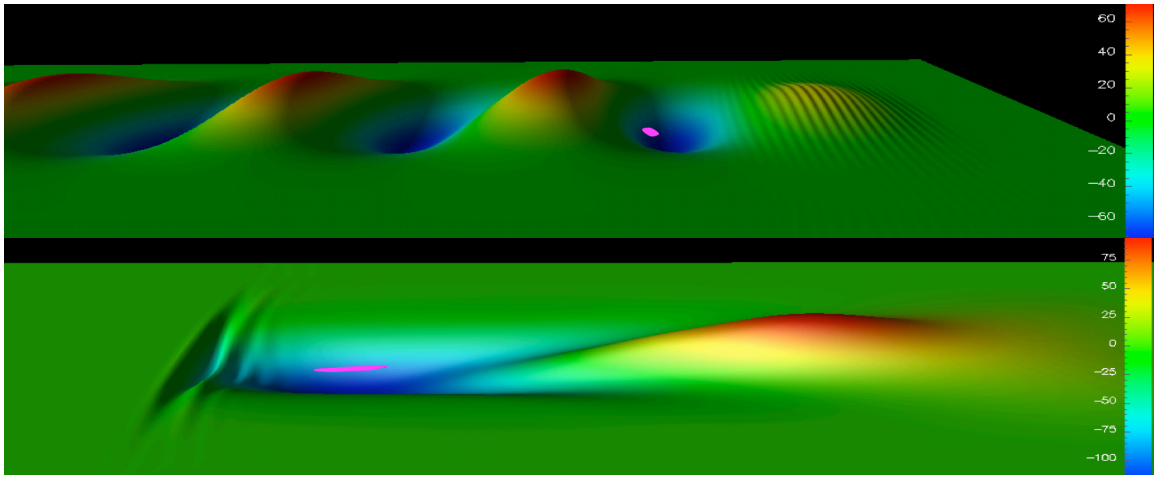

Fig. 26: Rendering from a 2-1/2D Warp simulation of a laser-plasma wakefield accelerator stage in the laboratory frame (top) and a boosted frame at γ=10 (bottom), showing a colored surface plot of the longitudinal electric field exhibiting the imprint from the driving laser pulse, and the plasma wake. The accelerated electron beam (magenta) is loaded in the first period of the wake.

Warp simulations of a 10 GeV class stage of laser wakefield accelerator at plasma density $n_e=10^{19}$ cm^{-3} were performed in 2-1/2D and 3D using reference frames moving anywhere between γ=1 (laboratory frame) and 13. These simulations are scaled replicas of 10 GeV stages that would operate at actual densities of 10^{17} cm^{-3} [23,24] and allow short run times to permit effective benchmarking between the algorithms. Agreement

within a few percent was observed on the beam peak energy and average energy between calculations in all frames, showing that the boosted frame simulations gain speed without sacrificing accuracy. Speedups of 200 and 130 were demonstrated in 2-1/2D and 3D respectively between the calculation in the frame at γ=13 and the calculation in the laboratory frame. For full scale simulations, speedups over 100,000 in 3D and 1 million in 2-1/2D were demonstrated for 100 GeV and 1 TeV stages respectively [12].

The average beam energy as a function of position in the laboratory frame and the CPU time recorded as the beam crosses successive stations are plotted in Figures 27 for scaled runs in a frame of reference boosted at γ=13.

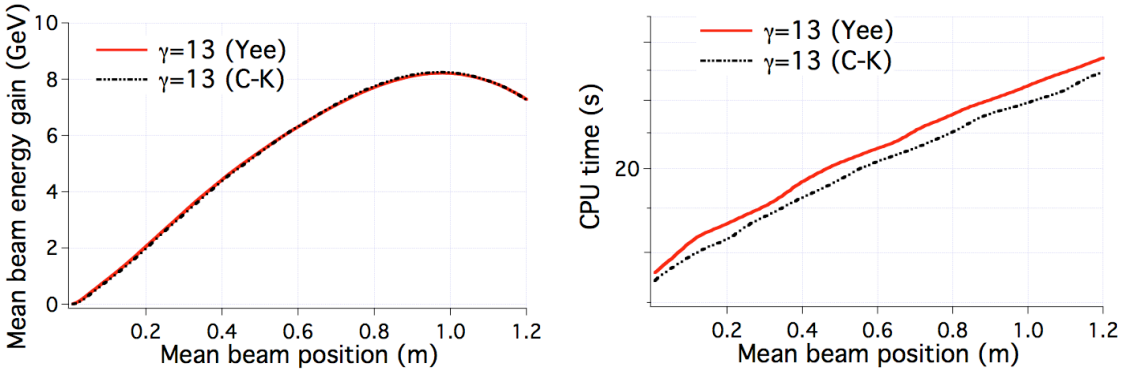

Fig. 27: (left) average beam energy versus longitudinal position in the laboratory frame from 2-1/2D Warp simulations in a frame moving at γ=13; (right) CPU time recorded as the beam crosses successive stations in the laboratory frame, using the Yee or the Cole-Karkkainen solver.

Excellent agreement is observed between calculations using the Yee solver and the Cole-Karkkainen solver, validating further Warp's PIC implementation of the CK solver. Although the CK solver has a larger stencil and is thus more computationally expensive, it allows for a larger time step than the Yee solver, resulting in a slightly shorter run time.

Mesh refinement was applied for reducing the number of mesh cells and macroparticles of a 2-1/2D simulation in a boosted frame (γ=10) of a scaled 10 GeV class stage. For this first test, up to three levels of mesh refinement were used (see Fig. 28 – top), each level refining by a factor of two from the previous level (in the transverse direction only). For these intial tests, no adaptation of macroparticles weight and/or number was performed, and the same time step verifying the Courant condition (using the Yee solver) at the finest lever was used for all levels. The beam transverse emittance is plotted in Fig. 28 as a function of distance in the plasma for (i) four runs using a unique grid (no mesh refinement) with respectively 100, 200, 400 and 800 cells in the transverse direction; (ii) three runs using a main grid with 100 cells transversely and respectively one, two and three levels of mesh refinement. The results obtained with a unique grid at a given resolution are well recovered by using a coarse main grid and the appropriate number of refinement levels to match the resolution of the single grid runs. This shows that the mesh refinement algorithm implemented in Warp allows relaxation of the

resolution of the main grid transversely while preserving the overall accuracy as measured on the evolution of the beam emittance.

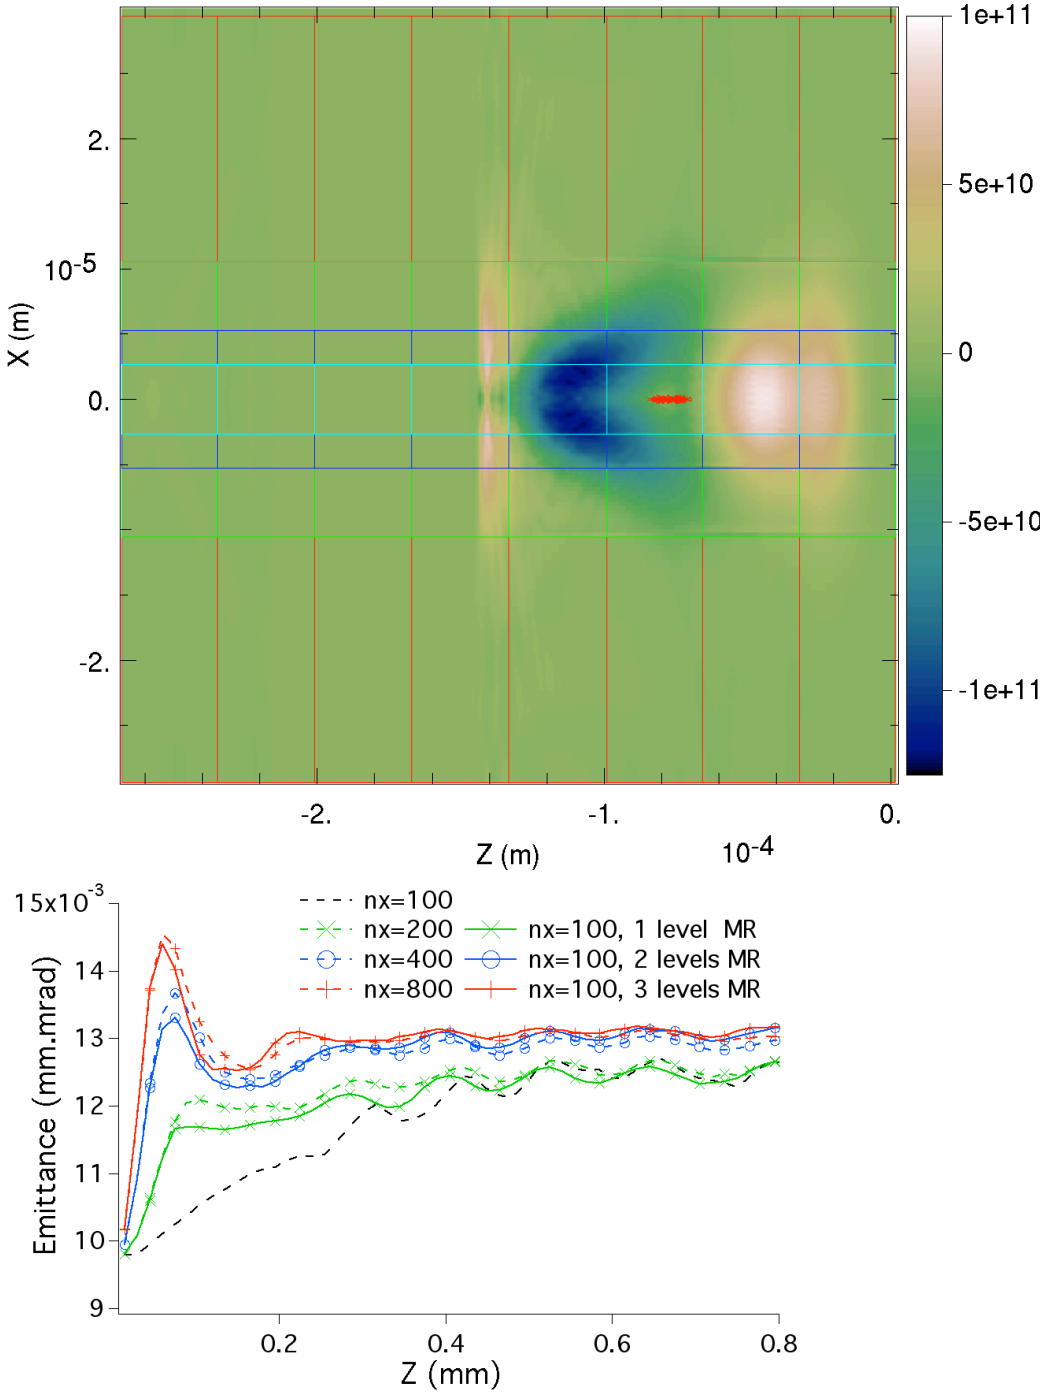

Fig. 28: (top) color plot of the longitudinal electric field magnitude from a simulation of a scaled 10 GeV LPA stage using a Lorentz boosted frame of reference with γ=10, with 3 level of refinements delimited by red, green, blue and cyan boxes (rendered from a 8 processors run with 1D longitudinal domain decomposition, each processor rendering its own domain); (bottom) beam transverse emittance as a function of distance in the plasma for (i) four runs using a unique grid (no mesh refinement) with respectively 100, 200, 400

and 800 cells in the transverse direction; (ii) three runs using a main grid with 100 cells transversely and respectively one, two and three levels of mesh refinement

3 - Conclusion

A novel electromagnetic solver with mesh refinement capability was successfully implemented in Warp. The solver allows for calculations in 2-1/2 and 3 dimensions, includes the standard Yee stencil, and the Cole-Karkkainen stencil for lower numerical dispersion along the principal axes.

Warp implementation of the Cole-Karkkainen stencil includes an extension to perfectly matched layers (PML) for absorption of waves, and is preserving the conservation property of charge conserving current deposition schemes, like the Buneman-Villanesor and Esirkepov methods.

Warp's mesh refinement framework (originally developed for electrostatic calculations) was augmented to allow for electromagnetic capability, following the methodology presented in [1] extended to an arbitrary number of refinement levels.

Other developments include a generalized particle injection method, internal conductors using stair-cased approximation, and subcycling of particle pushing. The solver runs in parallel using MPI message passing, with a choice at runtime of 1D, 2D and 3D domain decomposition, and is shown to scale linearly on a test problem up-to 32,768 CPUs.

The novel solver was tested on the modeling of filamentation instability, fast ignition, ion beam induced plasma wake, and laser plasma acceleration, on which it was shown that Warp's mesh refinement capability could successfully reproduce calculations using a single high resolution grid at a fraction of the cost.

Further work is needed for fully integrating the various pieces (mesh refinement, internal conductors, particle emission, particle subcycling) in a fully coherent framework, as well as to explore adaptivity, where refinement patches are created and removed dynamically.

References

[1] J.-L. Vay, J.-C. Adam, and A. Héron, "Asymmetric PML for the absorption of waves. Application to mesh refinement in electromagnetic particle-in-cell plasma simulations," *Comp.Phys. Comm.* **164** (2004) 171-177.

[2] J.-P. Bérenger, "A perfectly matched layer for the absorption of electromagnetic waves," *J. Comput. Phys.* **114** (1994) 185–200.

[3] J. B. Cole, " A high-accuracy realization of the Yee algorithm using non-standard finite differences", *IEEE Trans. Microw. Theory Tech.*, **45** (1997) 991–996.

[4] J. B. Cole, "A high-accuracy Yee algorithm based on nonstandard finite differences: new developments and verifications", *IEEE Trans. Antennas Prop.*, **50** (2002) 1185–1191.

[5] M. Karkkainen, E. Gjonaj, T. Lau, T. Weiland, "Low-Dispersion Wake Field Calculation Tools", Proceedings of ICAP 2006, Chamonix, France.

[6] C. K. Birdsall and A. B. Langdon, Plasma Physics Via Computer Simulation (Adam-

Hilger, 1991).

[7] A. B. Langdon, *Comput. Phys. Comm.* **70** (1992) 447.

[8] B. Marder*, J. Comput. Phys.* **68** (1987) 48.

[9] J.-L. Vay, C. Deutsch, *Phys. Plasmas* **5** (1998) 1190.

[10] J. Villasenor, O. Buneman, *Comput. Phys. Comm.* **69** (1992) 306.

[11] T. Zh. Esirkepov, "Exact charge conservation scheme for Particle-in-Cell simulation with an arbitrary form-factor," *Comp. Phys. Comm.* **135** (2001) 144.

[12] J.-L. Vay, E. Cormier-Michel, C. G. R. Geddes, D. P. Grote, "Modeling laser wakefield accelerators in a Lorentz boosted frame.", *in preparation*

[13] J.-L. Vay, *J. Comput. Phys.* **165** (2000) 511.

[14] J.-L. Vay, "Asymmetric perfectly matched layer for the absorption of waves," *J. of Comp. Phys.*, **183** (2002) 367–399.

[15] S. A. Cummer, "A simple, nearly perfectly matched layer for general electromagnetic media," *IEEE Microwave Wireless Components Lett*., **13** (2003) 128–130.

[16] W. Hu and S. A. Cummer, "The nearly perfectly matched layer is a perfectly matched layer," *IEEE Antennas And Wireless Propagation Lett*., **3** (2004) 137-140.

[17] J.-L. Vay and C. Deutsch, "A three-dimensional electromagnetic particle-in-cell code to simulate heavy ion beam propagation in the reaction chamber," *Fus. Eng. And Design* **32-33** (1996) 467.

[18] J. J. Watrous, J. W. Luginsland, G. E. Sasser III, "An improved space-charge-limited emission algorithm for use in particle-in-cell codes", Phys. Plasmas 8, 289 (2001); DOI:10.1063/1.1329152.

[19] J.-C. Adam, A. Gourdin Serveniere, A. B. Langdon, "Sub-Cycling in Particle Simulation of plasma", *Journal of Comput. Physics*, **47** (1982) 229-244.

[20] A. Friedman, S. E. Parker, S. L. Ray, and C. K. Birdsall, "Multi-scale Particle-in-Cell Plasma Simulation", *J. Comput. Phys.* **96** (1991) 54.

[21] I. D. Kaganovitch, *et al.,* "Nonlinear plasma waves excitation by intense ion beams in background plasma," *Phys. Plasmas* **11** (2004) 3546.

[22] C.G.R. Geddes et al, 'Computational studies and optimization of wakefield accelerators', J. Phys. Conf. Series V **125** (2008) 12002/1-11.

[23] E. Cormier-Michel, *et al.*, *Proc. 13th Advanced Accelerator Concepts Workshop*}, Santa Cruz, CA (2008) 297.

[24] C. G. R. Geddes *et al*., *Proc. Particle Accelerator Conference*, Vancouver, Canada (2009) WE6RFP075.

[25] J.-L. Vay, 'Noninvariance of space- and time- scale ranges under a Lorentz transformation and the implications for the study of relativistic interactions', *Phys. Rev. Lett.* **98** (2007) 130405.