Application of adaptive mesh refinement to particle-in-cell simulations of plasmas and beams[a]

J.-L. Vay,[b] P. Colella, J. W. Kwan, P. McCorquodale, and D. B. Serafini
Lawrence Berkeley National Laboratory, Berkeley, California 94720

A. Friedman, D. P. Grote, and G. Westenskow
Lawrence Livermore National Laboratory, Livermore, California 94550

J.-C. Adam and A. Héron
CPHT, Ecole Polytechnique, 91128 Palaiseau, France

I. Haber
University of Maryland, College Park, Maryland 20742-3511

Plasma simulations are often rendered challenging by the disparity of scales in time and in space which must be resolved. When these disparities are in distinctive zones of the simulation domain, a method which has proven to be effective in other areas (e.g., fluid dynamics simulations) is the mesh refinement technique. A brief discussion of the challenges posed by coupling this technique with plasma particle-in-cell simulations is given, followed by a presentation of examples of application in heavy ion fusion and related fields which illustrate the effectiveness of the approach. Finally, a report is given on the status of a collaboration under way at Lawrence Berkeley National Laboratory between the Applied Numerical Algorithms Group (ANAG) and the Heavy Ion Fusion group to upgrade ANAG's mesh refinement library Chombo to include the tools needed by particle-in-cell simulation codes. © *2004 American Institute of Physics.*
[DOI: 10.1063/1.1689669]

I. INTRODUCTION

Plasma simulations are often rendered challenging by the disparity of scales in time and in space which must be resolved. This is true, for example, for the modeling in detail of the beam behavior from end-to-end in a heavy ion fusion (HIF) (see, for example, Ref. 1) accelerator. Although such a simulation is ultimately needed, it is at present beyond reach, mostly due to the fact that the space scales span up to nine orders of magnitude, from microns (the Debye length close to emitter) to a kilometer or more (the length of the accelerator).

The mesh refinement (MR) technique addresses the issue of a wide range of space scales and has been applied successfully in other fields (i.e., fluid dynamics). While preserving most of the advantages of regular gridding, this technique covers areas that need a higher resolution with a finer mesh. If the areas of the physical domain that need refinement evolve in time, then an automatic redistribution of the refinement following some criterion applies: this is known as adaptive mesh refinement (AMR).

Section II provides a brief discussion of the challenges posed by coupling this technique with plasma particle-in-cell (PIC) simulations, for both the electrostatic (Vlasov–Poisson equation set, labeled "AMR/ES–PIC") and the electromagnetic (Vlasov–Maxwell equation set, labeled "AMR/EM–PIC") cases. In Sec. III, we present examples of AMR/ES–

PIC in application to the modeling of HIF ion sources, in steady-state and in time-dependent modes, while an application of a new AMR/EM–PIC method to the modeling of laser-plasma interaction in the context of fast ignition is given in Sec. IV. Section V reports on the status of a collaboration under way at Lawrence Berkeley National Laboratory (LBNL) between the Applied Numerical Algorithms Group (ANAG) and the Heavy Ion Fusion group to upgrade ANAG's mesh refinement library Chombo to include the tools needed by particle-in-cell simulation codes. Conclusion and further comments are offered in Sec. VI.

II. MESH REFINEMENT FOR PARTICLE-IN-CELL SIMULATIONS

The application of the mesh refinement technique to particle-in-cell simulations requires special care. Using prototypes, we have performed studies of various issues. We summarize here our findings and refer the reader to Ref. 2 for a more detailed presentation.

The use of MR (or AMR) implies a breaking in symmetry in the field solution which may translate into a spurious self-force on any macroparticle traveling close to a refinement patch boundary. This is a generic problem to any implementation of MR in PIC (in fact, to any particle-mesh code using irregular gridding). Depending on the implementation chosen to introduce MR into PIC, there are various ways to mitigate this effect. In addition, some implementations may introduce a violation of Gauss' law at the discrete level, which means that the total charge may not be con-

---

served. Although both of these effects are expected to be of small instantaneous magnitude in most cases, their consequences in simulations involving many time steps, where significant accumulation of errors may occur, are just beginning to be explored. These effects must be taken into account when making the choice of the implementation.

The application of mesh refinement to electromagnetic simulations encounters an additional difficulty specific to the modeling of the wave equation on a nonuniform mesh. A wavelength which is resolved on a finely gridded region of the system may not be described on a coarser gridded region, and is reflected at the interface. As shown in Ref. 3, these spurious reflected waves have a magnitude larger than the original wave with many mesh refinement schemes, leading eventually to an instability due to multiple reflections. These spurious reflections can be either damped inside the entire grid using an algorithm like the one described in Ref. 4, or absorbed at the fine-to-coarse interface using a specialized algorithm (see Refs. 3 and 5). We have identified two approaches. In the first, a complex algorithm is used to ''sew'' a fine gridded patch into a hole in a parent coarser grid, as proposed in Ref. 3 (for example). In the second, as proposed in Ref. 5, we solve the system entirely on the coarse grid, then compute a correction by solving on a set of two patches limited by absorbing boundary conditions, once at each resolution. When gathering the field onto the macroparticles, we substract the standard resolution solution from the global solution and add the fine grid solution, in effect substituting the coarse grid solution by the fine grid one in the refined region. The second method offers greater simplicity of implementation and removes the problem of potential instability due to reflected waves at the fine-to-coarse interface. However, it is computationally slightly more costly, and it may introduce spurious waves when electromagnetic solvers with large numerical dispersion are used inside the patch (see Ref. 5). Clearly, each method has its advantages and its drawbacks, and both are under investigation. Since the first method has already been partially analyzed and tested in Ref. 3, we concentrate in this paper on the second approach, which is new and was in need of testing.

Despite the difficulties, mesh refinement can be used effectively in particle-in-cell simulations, as shown by the following example applications performed with prototypes of mesh refinement methods implemented in existing PIC codes.

III. APPLICATION OF AMR/ES–PIC TO ION SOURCE MODELING

Figure 1 shows a snapshot taken from a movie of an end-to-end simulation of the High-Current Experiment (HCX)[6] at LBNL, using HIF program's accelerator code WARP (multidimensional PIC electrostatic, see Ref. 7). It shows the beam emitted from a triode source and traveling through the four accelerating and focusing electrostatic quadrupoles of the injector. The modeling of the triode region is critical since it determines the initial phase-space shape of the beam. It is also a challenging section to model since there is a large variation in particle density close to the emitter, and
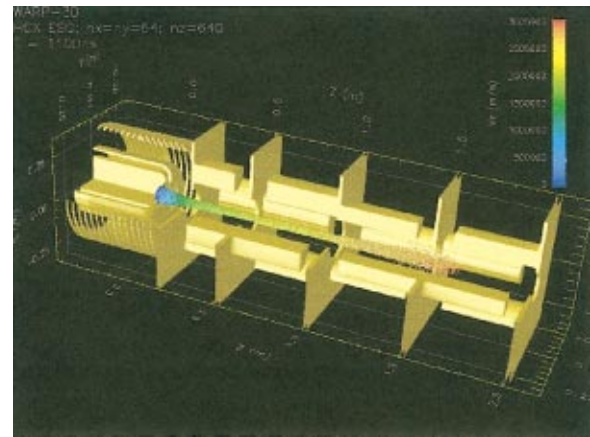

FIG. 1. (Color) Three-dimensional rendering of HCX injector simulation from a movie of an end-to-end WARP simulation of the HCX experiment. This shows the beam, emitted from the source (left), propagating through the first quadrupole lenses.

an accurate description of the emitter and beam edges is crucial. The run shown in the figure was carried out before AMR–PIC was developed and so is not ''converged.'' Nonetheless, the results were usable for the purposes intended, and the figure is indicative of the kind of problem to be solved.

A. Steady state

We display in Fig. 2 (top) a snapshot of the beam taken from a quasisteady state axisymmetric WARP simulation of the triode. Here, ''quasisteady state'' means a time-dependent calculation of the beam being emitted from the source, solving for the field every n time steps with $10 < n < 50$ typically. The simulation is stopped when an equilibrium is reached, under the assumption that the equilibrium solution exists and is unique (both of these assumptions are not guaranteed but have been found to be fulfilled in spot checking). A dashed line shows the limits of a fixed refinement patch surrounding the emitter (referred to later as ''MR''). Figure 2 (bottom) shows a different kind of grid structure that was used to
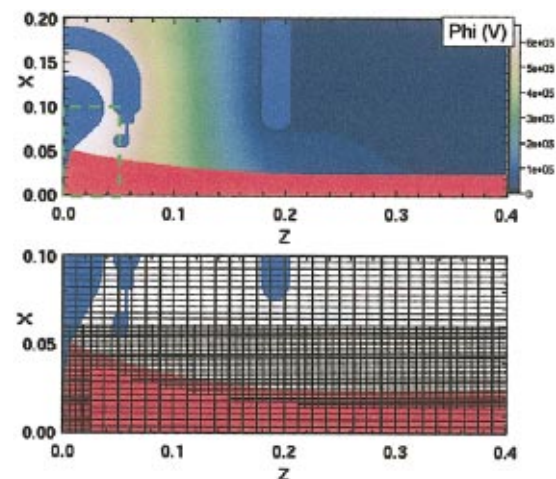

FIG. 2. (Color) Top: color contour plot of electric potential with triode structure (blue) and beam (red). Bottom: schematic of gridding when using AMR; the emitting area and the beam edge are covered with a finer grid.
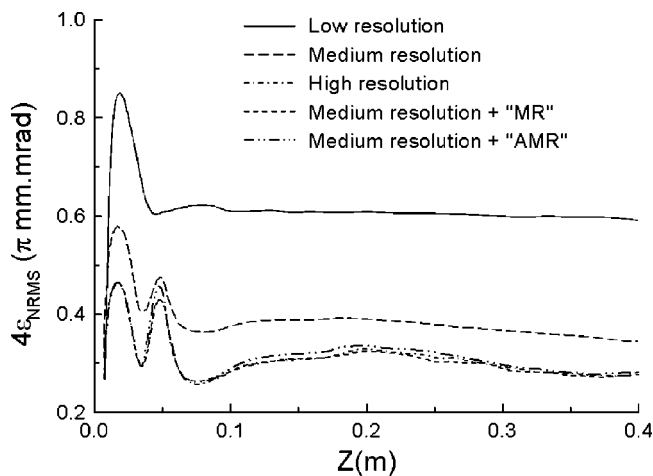
FIG. 3. Beam rms emittance (a figure of merit for beam quality—the lower the better) as a function of z (statistical noise was removed for clarity): the emittance converges downward with increasing resolution. The high resolution result is recovered with a run at medium resolution and the MR or AMR patch.

emulate an AMR set of patches. A finely gridded patch was tailored with a special shape so as to apply a fine gridding around the emitter and the beam edge, and a coarser gridding inside the beam. The shape of the grid was reset at each time step so as to follow the beam edge during the transient (referred to later as "AMR").

The root-mean-square (rms) normalized emittance versus Z (at steady-state) is shown in Fig. 3 for three runs using uniform low, medium and high resolution and two runs using medium resolution plus the MR or AMR refinement patch. The jump of resolution from low to medium and medium to high is a factor of 2 in each direction and a factor of 4 in the number of macroparticles, in order to keep the number of macroparticles per cell constant on average. The number of macroparticles used is the same in all the medium resolution runs, with or without MR or AMR. The "high resolution" case serves as a reference.

The results show that the emittance converges downward with increasing resolution and that the high resolution result is recovered from the medium resolution run with MR or AMR. Plots of phase-space projections and charge density at $Z=39$ cm (Fig. 4) confirm this and reveal more details: the size of the beam edge "hook" diminishes with increasing resolution, as does the spike at the edge ($R=2$ cm) of the charge density distribution. The size of the "hook" at high resolution is well recovered at medium resolution using either the MR or the AMR patch. However, the charge density spike is only well recovered using the AMR patch.

In this example, the use of a prototype AMR refinement patch has allowed recovery of essentially the same results as those of a reference case, at about a fourth of the computational cost. Higher gains are expected with the full three-dimensional production package Chombo (see later).

B. Time dependent

Figure 5 displays snapshots of the beam head in the injector of HCX for two different rise times of the voltage
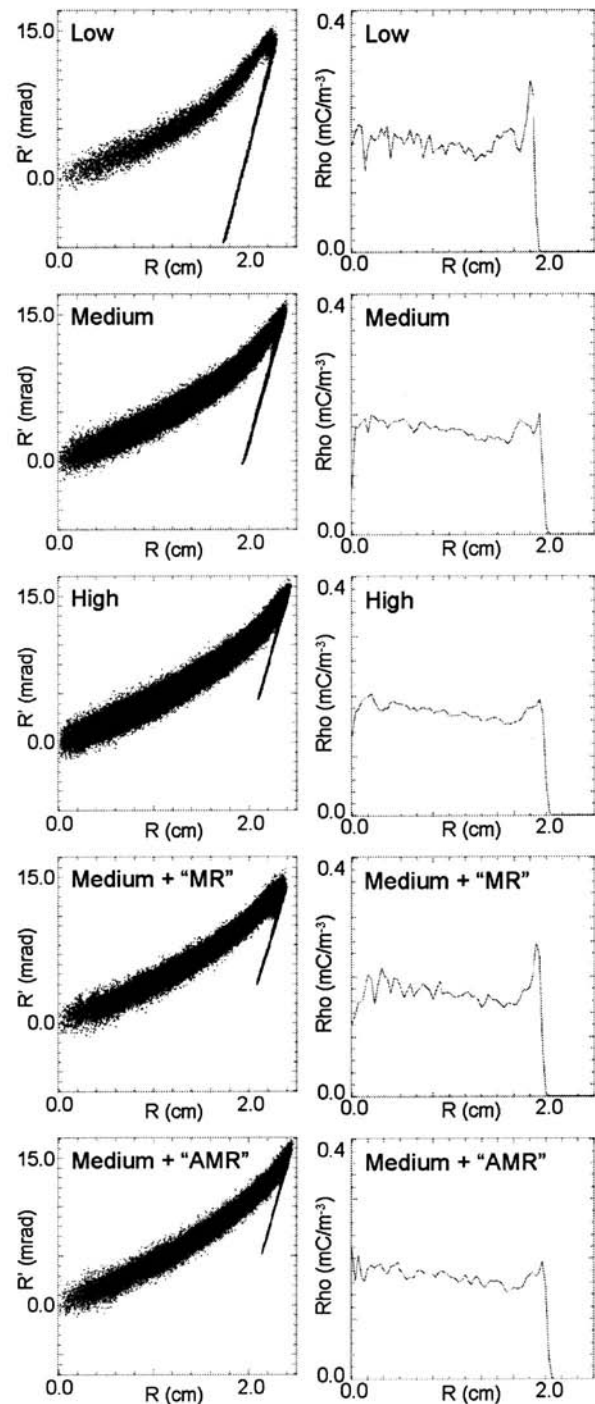

FIG. 4. Phase-space projection $R'=dR/dZ$ vs radius R and charge density versus R at $Z=0.39$ m for the cases with uniform grid at low, medium and high resolution, and medium resolution with MR or AMR patch. Part of the beam edge behavior, which was not adequately modeled using the MR patch covering the emitting area only, was captured using the AMR patch which dynamically refines the region surrounding the beam edge.

applied on the source gate, from a WARP simulation using the voltage profile measured in the experiment. With a rise time of 800 ns, the simulation predicts a significant "mismatch" of the head resulting in loss of a small fraction of particles at the wall. According to the second simulation, shortening the
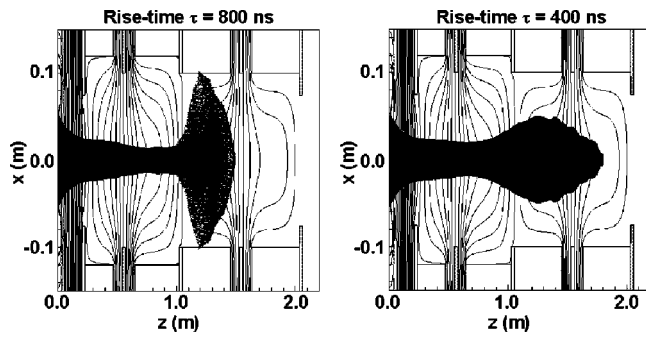
FIG. 5. Snapshot of the beam head in simulations of the HCX experiment for two different rise times of the applied voltage. On the left, the rise time was of 800 ns (wave form and rise time were taken from actual experimental data) and the simulation predicts an explosion of the beam head, leading to loss of particles at the wall. The same simulation using a shorter rise time (400 ns) predicts better control of the beam head.

rise time to 400 ns would provide a well-controlled beam head, keeping outer particle trajectories away from any structure.

The search for an optimum gate voltage profile has driven the use of mesh refinement, and development of a specialized version of it, for time-dependent modeling of the ion source. Since the core of the algorithm for injecting particles at the source is locally one-dimensional (with three-dimensional corrections where needed), we restricted ourselves to a one-dimensional study of an ideal voltage wave form first.

1. One-dimensional study

For a one-dimensional diode, Lampel and Tiefenback demonstrated that an analytical solution for a wave form of the applied voltage producing a Heaviside step for the current profile exists and is given by

$$V(t) = \frac{1}{3} \frac{t}{t_{transit}} \left[4 - \left(\frac{t}{t_{transit}} \right)^3 \right] V_{max}, \tag{1}$$

where V_{max} is the voltage that is applied at steady state and $t_{transit}$ is the transit time of a particle from the emitter to the collector (see Ref. 8). The latter quantity is given by

$$t_{transit} = 3 \cdot d \cdot \sqrt{\frac{m}{2qV_{max}}}, \tag{2}$$

where q and m represent the charge and mass of particles and d is the distance between the emitter and the collector. In Fig. 6, we show the Lampel–Tiefenback voltage profile as well as the current history obtained from three simulations.

The first simulation, using a single uniform grid, produced large unphysical oscillations damping with time. It was determined that these oscillations are due to insufficient resolution of the emission area and that the reduction of this oscillation to about a percent level required a scaling down of the mesh size by a factor of 10 000. In order to reach this goal, we developed a special version of an irregular gridded refinement patch, for which the gridding is set so that the number of macroparticles per mesh is a constant once the diode has reached steady-state Child–Langmuir flow. As shown in Fig. 6 (bottom left), the use of such a refinement
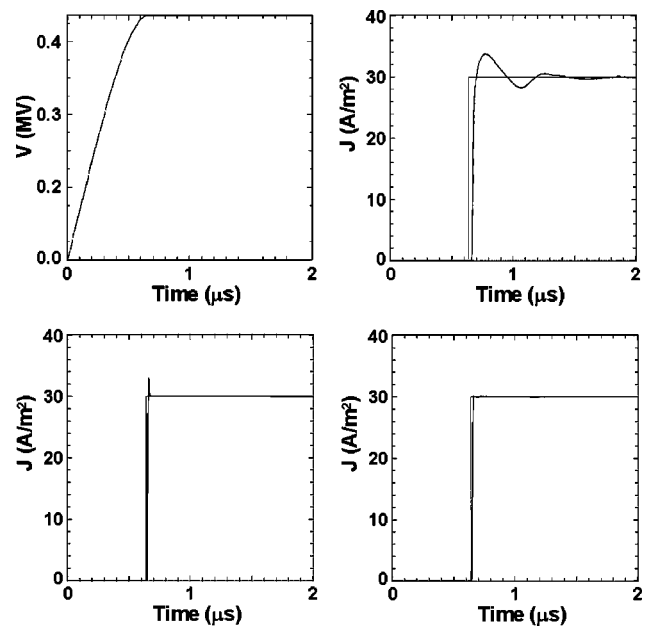

FIG. 6. Voltage profile given by the Lampel–Tiefenback solution for a Heaviside step current history in one dimension (top left). Comparison of current history from the analytic solution (Heaviside step) and from simulation: without mesh refinement (top right), with special mesh refinement patch around emitter (bottom left), and with an additional AMR patch following the beam head (bottom right).

patch effectively eliminates the spurious large oscillation. One spike remains, however, at the front of the current history; its origin was identified to be an insufficient resolution of the beam front as it propagates in the main grid. The last graph in Fig. 6 shows the current history from a simulation using the irregular specialized patch, as in the second case, and a regular adaptive refinement patch at 16 times the resolution of the main grid, following the front of the particle distribution. Both the spurious large amplitude oscillation and spike have been eliminated using the combination of fixed specialized and adaptive mesh refinement patches.

The specialized refinement patch developed in the one-dimensional case was implemented in the three-dimensional particle injection routines of WARP, as a two-dimensional array of one-dimensional patches, mapped onto the emitting surface (the direction of a line supporting a one-dimensional patch is perpendicular to the emitting surface).

2. Comparison of simulation versus experiment

Using a numerical technique introduced in Ref. 9 (a numerical three-dimensional extension of the analytical method used by Lampel and Tiefenback in one dimension), it is possible to derive from WARP the voltage profile corresponding to the emission of a very sharply rising current emitted from a source in a three-dimensional geometry. Such a profile was derived for the STS500 diode experiment (see Ref. 10) and is plotted in Fig. 7 together with the current predicted using three-dimensional WARP simulations (with axisymmetric representation of the field) with and without mesh refinement. Due to erosion of the beam head in the diode, the current profile history is not a heaviside step but is still very sharp.
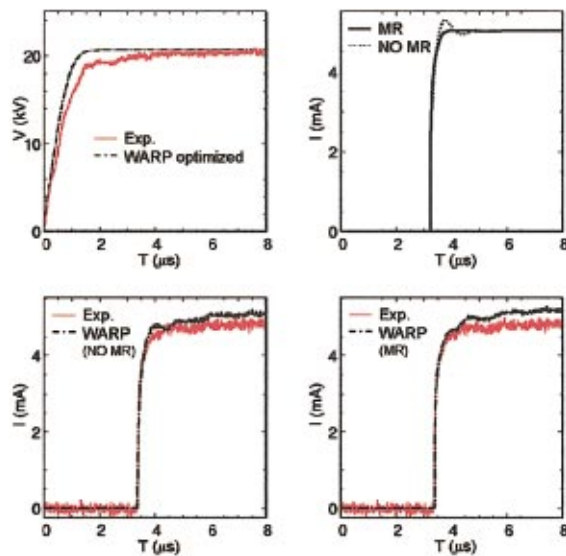
FIG. 7. (Color) Optimized voltage history predicted by WARP and as setup in the STS500 experiment (top left). Current history (statistical noise was removed for clarity) predicted by WARP using the optimized voltage profile, with and without mesh refinement (top right). WARP vs experiment current history using the experimental voltage profile, without (bottom left) and with (bottom right) mesh refinement. The current history obtained without using mesh refinement predicts an overshoot which is not present in the experimental history, validating the approach using mesh refinement.

Without mesh refinement, the simulation predicts an overshoot reminiscent of the large oscillation observed in the one-dimensional case.

   In order to check the simulation predictions against experiment, the voltage of the experiment was adjusted so as to match the predicted ''optimized'' voltage as closely as possible (Fig. 7). In practice, only the amplitude of the voltage could be modified; the shape of the voltage profile was not adjustable without hardware changes. Thus, the voltage was lowered in order to slow the beam to the point where the calculated transit time of a particle from emitter to exit of the diode matched the voltage rise time. This procedure gave the closest attainable experimental approximation of the optimized profile.

   This experimental profile was used in two WARP simulations, with and without mesh refinement (Fig. 7). The mesh refinement patch contained 50 cells, which translated into a factor of more than 1000 between the smallest cell size (closer to the emitter) and the size of the main grid cells. The early part of the experimental current profile is better matched by the simulation using the mesh refinement patch. The overshoot which was predicted by the simulation without mesh refinement is not present in the experimental result, as expected. A small disagreement exists between simulation and experiment on the value of the emitted current in steady state; within the experimental errors of the high voltage probe calibration. Future experiments will address this issue and will use a voltage profile closer to the optimized one, after a hardware upgrade.

IV. APPLICATION OF AMR/EM-PIC TO LASER–PLASMA INTERACTION MODELING

   The method presented in Ref. 5 has been implemented in the two-dimensional electromagnetic particle-in-cell code
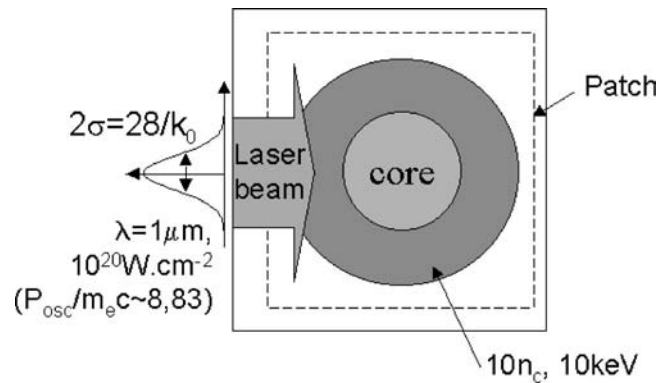

FIG. 8. Schematic of the test case. A laser impinges on a cylindrical target whose central part is artificially cooled to simulate a cold high-density core. The refinement patch encloses the entire target but the laser is launched outside the patch.

EMI2D, developed at Ecole Polytechnique for study of laser–plasma interaction in the context of fast ignition (see, for example, Ref. 11). In our test case (Fig. 8), a laser beam ($\lambda = 1 \ \mu m$, $I = 10^{20}$ W cm^{-2}, $2\sigma = 28/k_0$) impinges on a cylindrical plasma whose density ($10 n_c$) is far greater than the critical density (n_c) and whose initial temperature is 10 keV. The center of the plasma is artificially cooled to simulate a cold high-density core.

   Two cases have been run. A simulation with a single uniform high resolution grid is the case of reference. The other case has one patch covering the entire plasma target at the same resolution as the grid in the reference case, while the surrounding space, including the region where the laser is launched into the simulation, is covered by a grid at lower resolution (a factor of 2 along each dimension). The number of particles is the same in each run. A selection of plots is displayed in Fig. 9, which contains snapshots at (or near) the end of the simulations of the $X-Y$ and $X-V_x$ ion phase-space projections, as well as the field B_z, for the reference and the test cases. We observe that the results are very similar in both cases. The only noticeable discrepancy is a small spurious wave (a residual of the incident laser wave) exiting the refinement patch. Its origin comes from the error in the numerical dispersion relation of the Yee field advance algorithm and is well understood (see Ref. 5 for more details). Remedies, including the use of less dispersive field advance algorithms, are under study. We note that the spurious wave is negligible in this case, since it is propagating in a region of little interest. No instability was observed, nor was there any wave reflection at the edge of the finely gridded patch. This preliminary test is thus very encouraging and has validated the method to the point where it will soon be used in production runs at Ecole Polytechnique.

V. DEVELOPMENT OF A MULTIDIMENSIONAL AMR POISSON SOLVER FOR PIC CODES IN CHOMBO

   A collaboration is under way at Lawrence Berkeley National Laboratory between the ANAG and the Heavy Ion
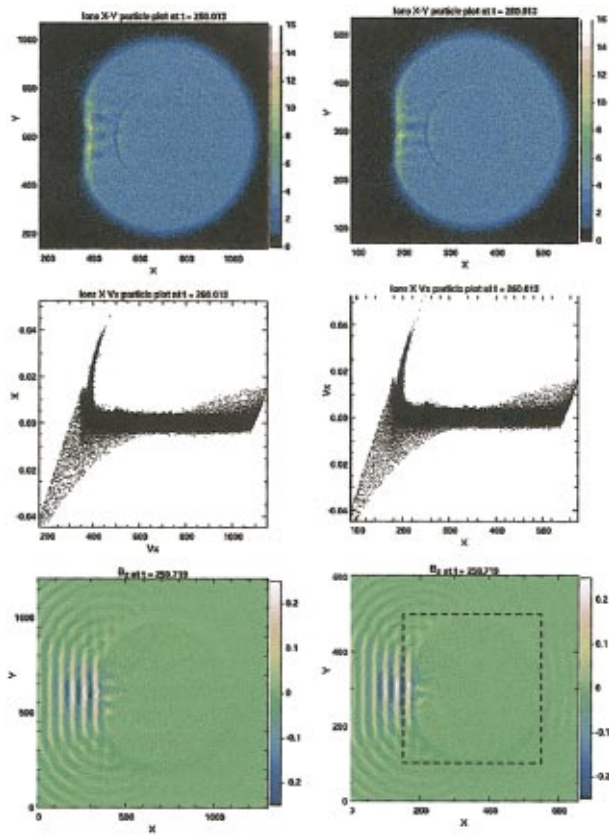
FIG. 9. (Color) $X-Y$, $X-V_x$ ion phase-space projections and B_z for the case of reference (left column) and the case with one refinement patch (right column). The limits of the patch are displayed as a dashed line on the plot of B_z.

Fusion group to upgrade ANAGs Chombo mesh refinement library with all the tools necessary for providing mesh refinement capability to existing PIC codes.

A nodal implementation of a multigrid AMR solver for the Poisson equation using Shortley–Weller (see Ref. 12) discretization of the Laplacian operator (sometimes referred as ''cut cell'' method), to account for internal boundaries (conductors) at subcell resolution, has been developed in Chombo (see Ref. 13). A prototype PIC interface to Chombo (called ChomboPIC), using this AMR Poisson solver, has been developed and integrated into the WARP code, mapping straightforwardly to the structure shown in Fig. 10. Although WARP is being used as the test PIC code during the development of this package, the latter is designed as an independent library which will ultimately be usable by any existing PIC code.

The three-dimensional solution of the electrostatic potential in the HCX injector was computed with Chombo us-
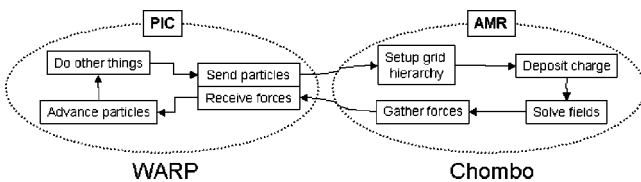

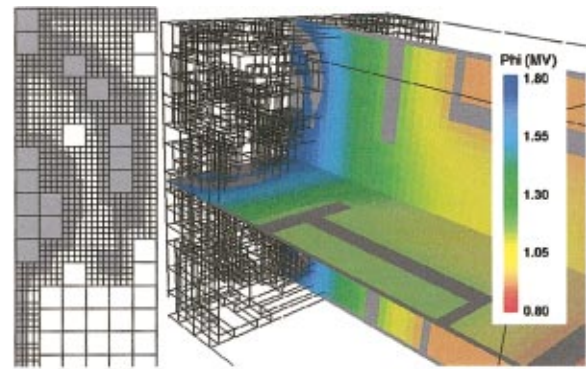FIG. 10. Diagram of WARP/Chombo configuration.

FIG. 11. (Color) Three-dimensional solution of the electrostatic potential in the HCX injector, as calculated by Chombo. A slice with the actual meshing (left) shows that the regions close to the boundaries of conductors (gray) are described with a finer mesh. The picture on the right shows a three-dimensional rendering that includes two orthogonal slices of the solution (with magnitude of electrostatic potential shown colored, conductors in gray) and the edges of the different domains containing finer mesh spacing (in this case, mesh refinement covered the conductor edges only in the area surrounding the source).

ing its automatic meshing capability. The criterion for refinement was to refine volumes covering the edges of the conductors of the source ($z<0.1$ m), with a ratio between coarse and fine mesh of 4. The result is displayed in Fig. 11 and shows how Chombo handles a complicated structure of grid blocks to get to the required solution nearly optimally. The criterion used for refinement was chosen to demonstrate the capability. Different criteria may be devised for actual modeling of the injector, for example refining the emitting region and the beam edge, as was successfully used in the HCX source example described earlier.

VI. CONCLUSION

After a brief discussion of the issues linked to the use of mesh refinement in particle-in-cell simulations, we have presented examples of applications in heavy ion fusion and laser-driven fast ignition using prototype implementations of mesh refinement algorithms in the HIF electrostatic PIC code WARP and the electromagnetic PIC code EMI2D from Ecole Polytechnique. In electrostatic PIC simulations, the results have demonstrated savings in computer memory and time by a factor of 4 for a typical scenario, and a factor of 1000 on a specific problem. Electromagnetic PIC simulations with mesh refinement are more challenging and the result presented here was limited to a proof of principle. Production runs which are intended to demonstrate significant savings with the use of mesh refinement are planned and will be reported when available.

In each of the examples presented here, the same time step was used to advance all the particles, regardless of the simulated region they were in. However, multiscale plasma simulations call for different time scales in separated regions of the simulated domain, using for example the technique developed in Ref. 14. This was not considered here in order to simplify both the implementations and the analysis. This aspect will be addressed in future work.

A collaboration to provide a library of mesh refinement tools for particle-in-cell codes is under way at LBNL, restricted to electrostatic solvers for now, but extending to electromagnetic solvers in the future. Typical gains of a factor of 10 or more are expected with Chombo once discrete-particle support is fully integrated.

ACKNOWLEDGMENTS

[1] C. M. Celata, F. M. Bieniosek, E. Henestroza *et al.*, Phys. Plasmas **10**, 2064 (2003).

[2] J.-L. Vay, P. Colella, P. McCorquodale, B. Van Straalen, A. Friedman, and D. P. Grote, Laser Part. Beams **20**, 569 (2002).

[3] J.-L. Vay, J. Comput. Phys. **167**, 72 (2001).

[4] A. Friedman, J. Comput. Phys. **90**, 292 (1990).

[5] J.-L. Vay, J.-C. Adam, and A. Héron, Comput. Phys. Commun. (to be published).

[6] P. A. Seidl, D. Baca, F. M. Bieniosek, A. Faltens, S. M. Lund, A. W. Molvik, L. R. Prost, and W. L. Waldron, Laser Part. Beams **20**, 435 (2002).

[7] D. P. Grote, A. Friedman, I. Haber, W. Fawley, and J.-L. Vay, Nucl. Instrum. Methods Phys. Res. A **415**, 428 (1998).

[8] M. Lampel and M. Tiefenback, Appl. Phys. Lett. **43**, 1 (1983).

[9] J.-L. Vay, A. Friedman, and D. P. Grote, *Proceedings of 7th International Computational Accelerator Physics Conference, East Lansing* (IOP, Bristol, 2002).

[10] J. Kwan, F. M. Bieniosek, J.-L. Vay, E. Halaxa, G. Westenskow, and I. Haber, Particle Accelerator Conference, Portland, 2003, paper ID FPAB006.

[11] M. Tabak, J. Hammer, M. Eglinsky, W. L. Kruer, S. C. Wilks, J. Woodworth, E. M. Campbell, M. D. Perry, and R. J. Mason, Phys. Plasmas **1**, 1626 (1994).

[12] G. H. Shortley and R. Weller, J. Appl. Phys. **9**, 334 (1938).

[13] P. McCorquodale, P. Colella, D. P. Grote, and J.-L. Vay, J. Comput. Phys. (submitted).

[14] A. Friedman, S. E. Parker, S. L. Ray, and C. K. Birdsall, J. Comput. Phys. **96**, 54 (1991).