Laser and Particle Beams (2002), *20*, 569–575. Printed in the USA.
Copyright © 2002 Cambridge University Press 0263-0346/02 $16.00
DOI: 10.1017/S0263034602204139

Mesh refinement for particle-in-cell plasma simulations: Applications to and benefits for heavy ion fusion

J.-L. VAY,[1] P. COLELLA,[1] P. McCORQUODALE,[1] B. VAN STRAALEN,[1]
A. FRIEDMAN,[2] AND D.P. GROTE[2]

[1]Lawrence Berkeley National Laboratory, Berkeley, California 94704, USA
[2]Lawrence Livermore National Laboratory, Livermore, California 94551, USA

(RECEIVED 27 May 2002; ACCEPTED 17 June 2002)

Abstract

The numerical simulation of the driving beams in a heavy ion fusion power plant is a challenging task, and simulation of the power plant as a whole, or even of the driver, is not yet possible. Despite the rapid progress in computer power, past and anticipated, one must consider the use of the most advanced numerical techniques, if we are to reach our goal expeditiously. One of the difficulties of these simulations resides in the disparity of scales, in time and in space, which must be resolved. When these disparities are in distinctive zones of the simulation region, a method which has proven to be effective in other areas (e.g., fluid dynamics simulations) is the mesh refinement technique. We discuss the challenges posed by the implementation of this technique into plasma simulations (due to the presence of particles and electromagnetic waves). We present the prospects for and projected benefits of its application to heavy ion fusion, in particular to the simulation of the ion source and the final beam propagation in the chamber. A collaboration project is under way at Lawrence Berkeley National Laboratory between the Applied Numerical Algorithms Group (ANAG) and the Heavy Ion Fusion group to couple the adaptive mesh refinement library CHOMBO developed by the ANAG group to the particle-in-cell accelerator code WARP developed by the Heavy Ion Fusion–Virtual National Laboratory. We describe our progress and present our initial findings.

Keywords: Heavy ion fusion; Mesh refinement; Particle in cell; Plasma simulation

1. INTRODUCTION

Integrated end-to-end simulation of a heavy ion fusion facility, or even the driver, is not yet possible with current software and hardware capabilities. Given the range of parameters that should be scanned by such types of simulations, relying solely on the progress in computer hardware power is not sufficient. Development and introduction into our codes of more efficient numerical techniques must be part of the effort. In this article, we report on the status of an ongoing effort aimed at introducing the adaptive mesh refinement (AMR) technique into particle-in-cell (PIC) modeling. This is part of a more general effort to apply AMR to plasma modeling. While we currently concentrate on implementing AMR into electrostatic PIC codes, we also consider their application to electromagnetic models and Vlasov simulations. In the first part of this article, we present the mesh refinement concept, a brief history of AMR (which was successfully developed for fluid modeling), and its potential benefits to heavy ion fusion. The second part will examine the main issues in applying AMR to models which contain particles and waves. The ongoing development of an electrostatic AMR–PIC code is based on the coupling of two existing packages: Chombo (http://seesar.lbl.gov/ANAG/chombo) for AMR and WARP (Grote *et al.*, 1996) for PIC. We present the approach and status of this effort in the third part, while results of AMR–PIC simulations realized with a RZ prototype developed in WARP are presented in the last part.

2. THE MESH REFINEMENT METHOD

The mesh refinement method (MR) is a technique for refining certain regions of the physical domain in a grid-based calculation. MR serves as a "numerical microscope," allowing researchers to "zoom in" on the specific regions of a problem that are most important to its solution. Rather than requiring that the whole calculation have the same spatial

Address correspondence and reprint requests to: Jean-Luc Vay, Mailstop 47-112, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94709, USA. E-mail: jlvay@lbl.gov

resolution, MR allows different resolutions in different regions of the problem. Areas of interest are covered with a finer mesh than the surrounding regions; for time-dependent problems, the finer meshes are also, in general, advanced with a smaller time step. The avoidance of having to perform the entire calculation at the finest resolution makes possible the solution of a wider range of problems. Sometimes, the areas of the physical domain that need refinement evolve and automatic redistribution of the refinement becomes necessary as the simulation unfolds; this is known as adaptive mesh refinement (AMR).

The hierarchical structured grid approach now known as AMR was first developed by Berger and Oliger (1984) for hyperbolic partial differential equations. The approach to adaptive gridding used here was developed for conservation laws and demonstrated to be highly successful for gas dynamics by Berger and Colella (1989) in two dimensions. Bell *et al.* (1991) extended the methodology to three dimensions. More recently, AMR has been extended to a variety of problems and algorithm choices, including, but not limited to, solving the variable-coefficient Poisson equation, the Helmholtz equation, the system of hyperbolic conservation laws governing inviscid gas dynamics, the compressible and incompressible Navier–Stokes equations, and the equations that govern reacting flows, such as those that occur in premixed and nonpremixed combustion. The heavy ion fusion program will benefit from the introduction of the AMR technique in several areas of its modeling effort, leading to integrated simulations of the driver, from source to target. The modeling of the particle injector in detail offers a specific challenge due to the presence of a singularity in the solution at the emitting surface. Convergence studies of the High-Current Experiment (HCX) injector at the Lawrence Berkeley National Laboratory (LBNL), using an axisymmetric (RZ) prototype AMR–Poisson solver, have shown that a very fine resolution is needed around the emitting surface (see below). In the accelerator, accurate modeling of halo generation calls for a detailed description of the beam edge which the AMR technique will render accessible through both mesh refinement and adaptivity capabilities. In the fusion chamber, a self-consistent simulation of a beam array is still out of reach, and the introduction of AMR would allow more rapid progress. To take full advantage of different approximations of the Vlasov–Maxwell system, we plan to introduce AMR into our PIC electrostatic and electromagnetic models, and ultimately into our Vlasov model. The implementation of the PIC electrostatic capability is under way, while the other implementations have been left for future work.

3. APPLICATION TO PARTICLE-IN-CELL PLASMA SIMULATION

The application of the mesh refinement technique to PIC plasma simulation needs special care to avoid the introduction of spurious effects into the model, or at least to minimize them. The main new challenge is the presence of macroparticles. The introduction of AMR introduces a spurious force that may potentially alter the particle motion to an unacceptable level. Also, some implementations of AMR may violate conservation laws and/or introduce nonphysical nonlinearities (anharmonic forces). Finally, in the case of electromagnetic PIC simulations, the introduction of AMR is challenging due to the reflection of high frequency waves at the boundary of a refinement patch and, in general, straightforward implementations of AMR into electromagnetic codes yield unstable algorithms.

3.1. Gauss theorem and field nonlinearities

Several methods can be envisioned to couple a fine grid and the coarser grid in which it is enclosed; we discuss two of them. The most straightforward method consists of, once the solution has been calculated to the desired precision on the coarse grid, getting the fine grid patch boundary values (for Dirichlet boundaries) through interpolation from the coarse grid solution. All the Dirichlet boundaries having been set on the fine grid, a solution is then computed in the interior, and the procedure is recursively performed for any refinement patch that it contains (this is the procedure that has been used in the example given in the next subsection "issues with macroparticles").

Another method, which is the default in the Chombo package (available at http://seesar.lbl.gov/ANAG/chombo), consists of iterating the solution back and forth between a patch and its mother grid. As in the other method, Dirichlet boundary values for the fine grid are interpolated from the coarse grid solution. Then, a specified number of iterations are performed in the fine grid and the fine and coarse grid solutions are reconciled during a "synchonization" step which consists of enforcing the fine grid solution on the coarse grid nodes located inside the fine grid patch. This procedure is iterated until convergence.

While the second method has been shown to be of higher order in accuracy, it violates Gauss' Law under a nodal implementation and modifies the coarse grid solution, eventually introducing otherwise nonpresent nonlinearities into that solution. These two effects may be issues for accelerator modeling. Both methods are being implemented and will be compared.

3.2. Issues with macroparticles

One cornerstone of PIC modeling consists of defining the charge deposition (from macroparticles to the grid) and field gathering (from the grid to the macroparticles) in such a way to avoid generating spurious self-force. However, it can be shown that, for electrostatic, the introduction of mesh refinement does create a spurious image of a particle at the interface between a coarse and a fine grid refinement patch. It is easy to demonstrate this effect on a simple example: a 64×64 square regular grid ("mother" grid) of undimen-

sioned length L with a local refinement patch consisting of a 32×32 square regular grid of length $L/4$. The centers of the mother grid and the patch grid coincide. A single particle is initialized at a position $\{X_0, Y_0\}$ in the fine grid. For each time step, the charge (undimensioned) of the particle is deposited on both grids using linear weighting. The field is computed on the coarse grid using a Poisson solver, assuming Dirichlet boundary condition $\varphi = 0$ at the boundaries. The solution is then linearly interpolated from the coarse grid onto the fine grid boundaries and the field solution is solved on the fine grid (the first method described in the preceding subsection). The field is then linearly interpolated back to the particle using the same weighting function as for the charge deposition. The particle is finally advanced using a leap-frog pusher. If the particle is outside the fine grid, only the coarse grid is considered for the calculation. Once the particle reaches the coarse grid boundaries, it experiences a specular reflection. The fine grid is reactivated as soon as the particle reenters it.

Figure 1 displays the results for a particle initialized at $\{X_0, Y_0\} = \{26, 32\}$ (in the coarse grid mesh size units) and $v_{x0} = v_{y0} = 0$. The results using the set of two grids as described above are compared to a reference case where the same problem is run on a uniform 128×128 square regular grid (so that it has the same resolution as the fine grid patch).

The particle is attracted by its image at the boundary and moves towards the closest boundary. As it should normally do and does in the reference case, the particle moves toward the closest boundary and bounces back, and so on. But in the

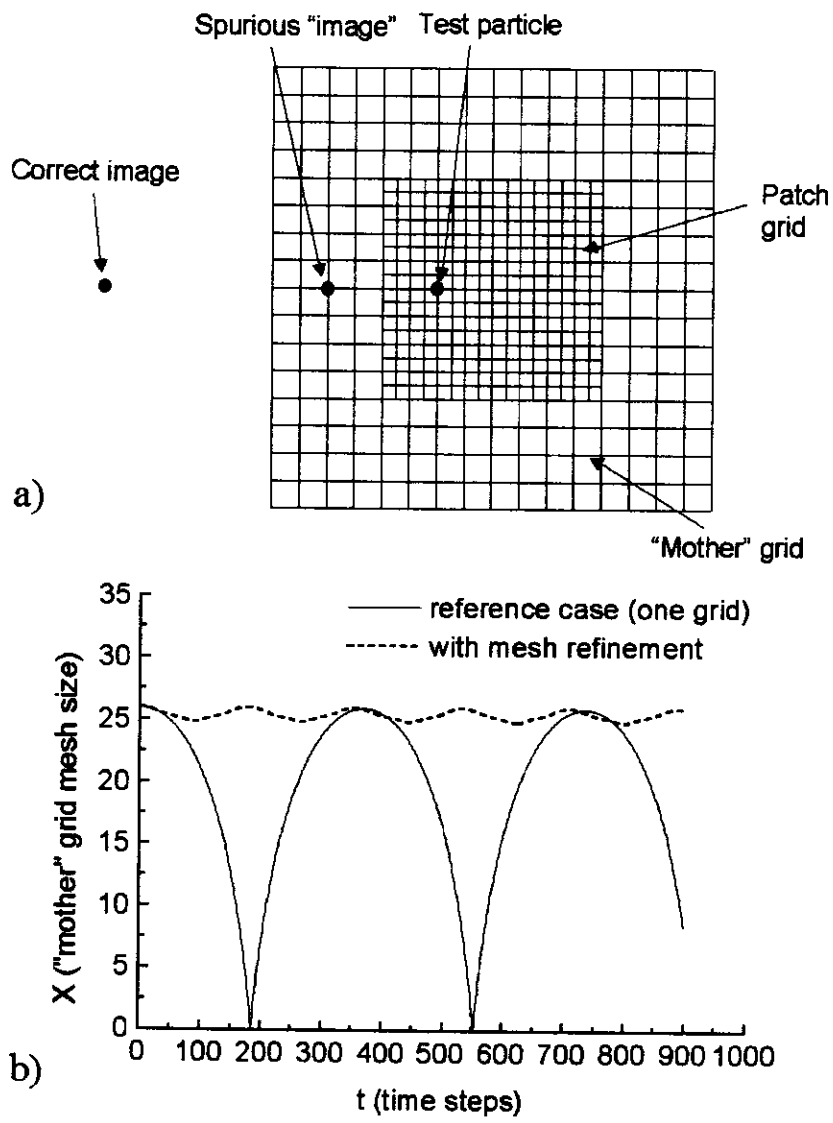

Fig. 1. Effect of mesh refinement on a single particle motion. Panel a shows the grids configuration. A test particle is initialized in the fine grid patch and is attracted by its image due to the metallic boundary (at the coarse grid border). In b, the time evolution of the particle position is given for an "exact" case (plain line) and for a case with mesh refinement (dashed line). The perturbation of the trajectory suggests that the equivalent of a spurious "image" of the real particle is introduced by the mesh refinement.

case of the two-grid set, the particle is reflected at the interface between the fine grid patch and the coarse mother grid, meaning that an effect equivalent to the presence of a spurious repulsive image force has been created at the interface by the local refinement. Techniques (not described here) to mitigate this effect have been explored and will be implemented if tests on real problems show that it is needed.

3.3. Electromagnetic waves

The modeling of electromagnetic waves poses a specific challenge to the use of mesh refinement and, more generally, to the use of non-Cartesian regular grids. This is due to the fact that wave propagation in vacuum is a process that has an "infinite memory" of past events so that errors tend to accumulate, leading to instability. By nature, a given grid can transport only waves whose wavelengths exceed roughly twice the mesh spacing. Hence, when refining an area by patching a fine grid on a coarse one, a band of wavelengths that can be resolved on the finer grid cannot be resolved on the coarse one, and are reflected at the fine-coarse grid interface. Such reflections are evidently unphysical and can build the field energy up by multiple reflections on the boundaries of the fine grid. Moreover, as shown in Vay (2001), the reflection of these waves at the interface is associated with an amplification of the reflected wave that can reach a factor of 10 in the case of refinement both in space and in time (smaller time steps in the refined region) of a factor of 3, for a scheme based on "jumps" of fine grid points over alternate coarse grid points and linear interpolations to connect the two grids. Some improvement can be obtained by applying an energy conserving scheme (Collino *et al.*, 1999) which forces, by construction, the reflection factor to be unity. Adjustable damping, as developed in Friedman (1990), can also help by damping the short wavelengths that constitute these spurious reflected waves, at the price of damping physical waves at the same frequencies. Another approach (Vay, 2001) consists in reducing the reflection factor by formally rewriting the multidimensional wave equation into monodimensional wave equations that are split upon the direction of propagation of the waves along an axis. Absorption of the considered waves at efficiency comparable to the first order, or eventually second order, Engquist and Majda (1977) absorbing boundary condition can be achieved using this technique. Further reduction of the reflection factor by orders of magnitudes may require the use of patches surrounded by perfectly-matched-layer (PML) (Berenger, 1994; J.-L. Vay, 2002) regions. For such a technique, for each refinement patch applied to a coarser grid, an additional patch having the same resolution as does the coarse grid will be needed. This technique relies on the fact that short wavelengths present in the fine-gridded patch will be created by sources located inside this same patch only. Hence, the short wavelengths can be computed on the fine patch, independently of the whole domain calculation, and be added to the coarse grid solution, pro-

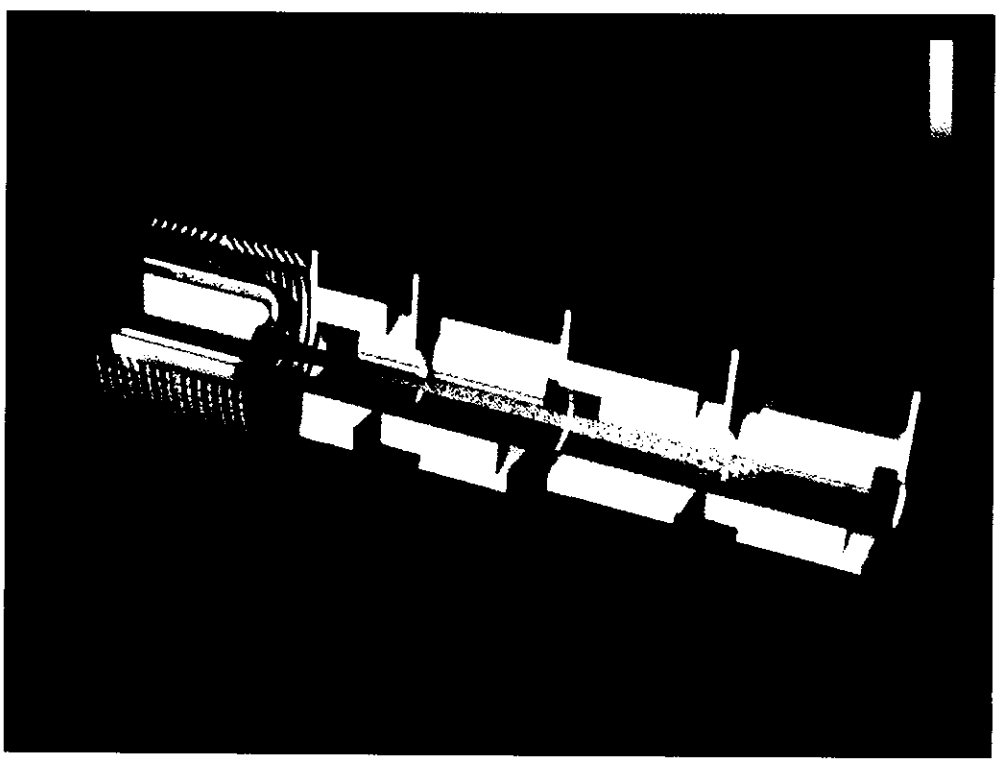

Fig. 2. Image from movie of an end-to-end WARP simulation of the HCX experiment. This shows the beam, emitted from the source (left), propagating through the first quadrupoles.

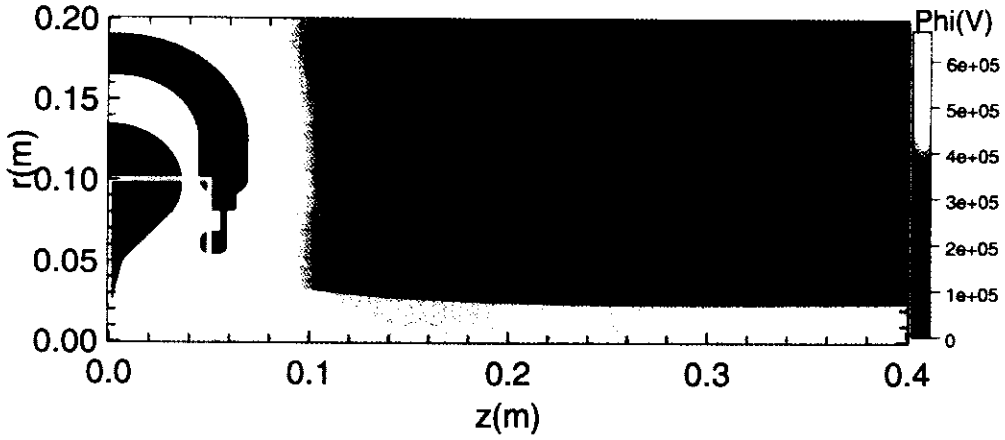

Fig. 3. Snapshot of the particle distribution (dark yellow) at the end of a HCX source simulation. Color contour plot of the electrostatic potential are rendered in vacuum areas (black areas represent conductors). The red box sets the limits of the mesh refinement patch that has been used in one of the runs.

vided that the long-wavelength information related to the same sources have also been calculated independently on the coarse grid patch and substracted from the whole domain solution. Compared to usual mesh refinement techniques, this approach has the advantage of removing the coupling between the coarse grid and the patch boundaries, removing associated problems and facilitating the implementation.

4. DEVELOPMENT OF THE WARP-CHOMBO PACKAGE

A nodal implementation of a multigrid AMR solver for the Poisson equation using Shortly–Weller ("cut cell") discretization of the Laplacian operator (to account for internal boundaries at subcell resolution) has been developed in Chombo. In our configuration, a library containing Chombo's

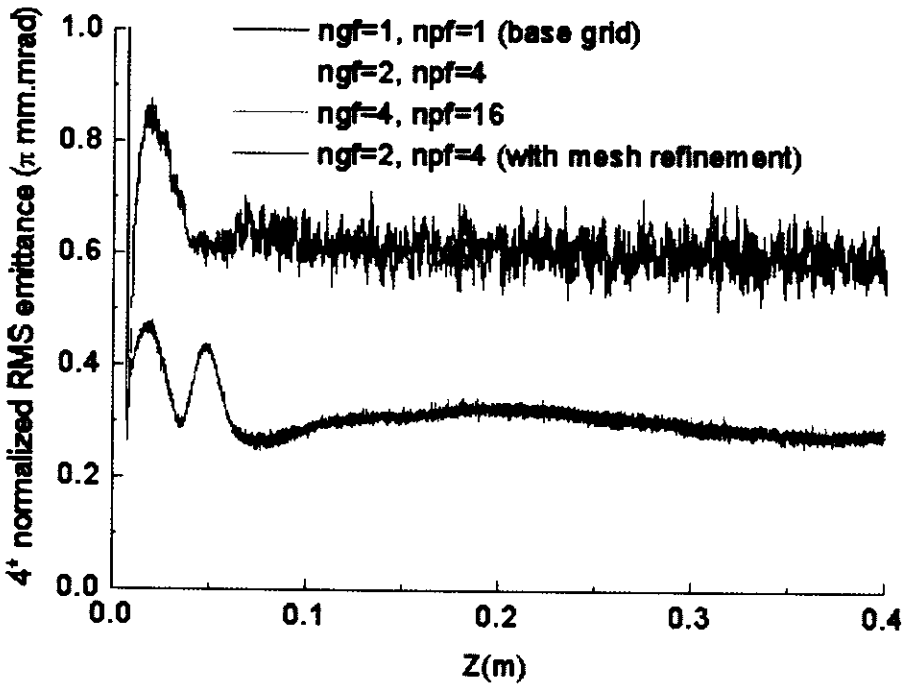

Fig. 4. Normalized rms emittance of the beam as a function of z, as measured at the end of four runs (see text for runs labeling description). A higher resolution leads to a lower emittance. Applying a mesh refinement around the emitting area allows to recover (almost) exactly the emittance profile at a reduce cost (factor of 4 in this case).

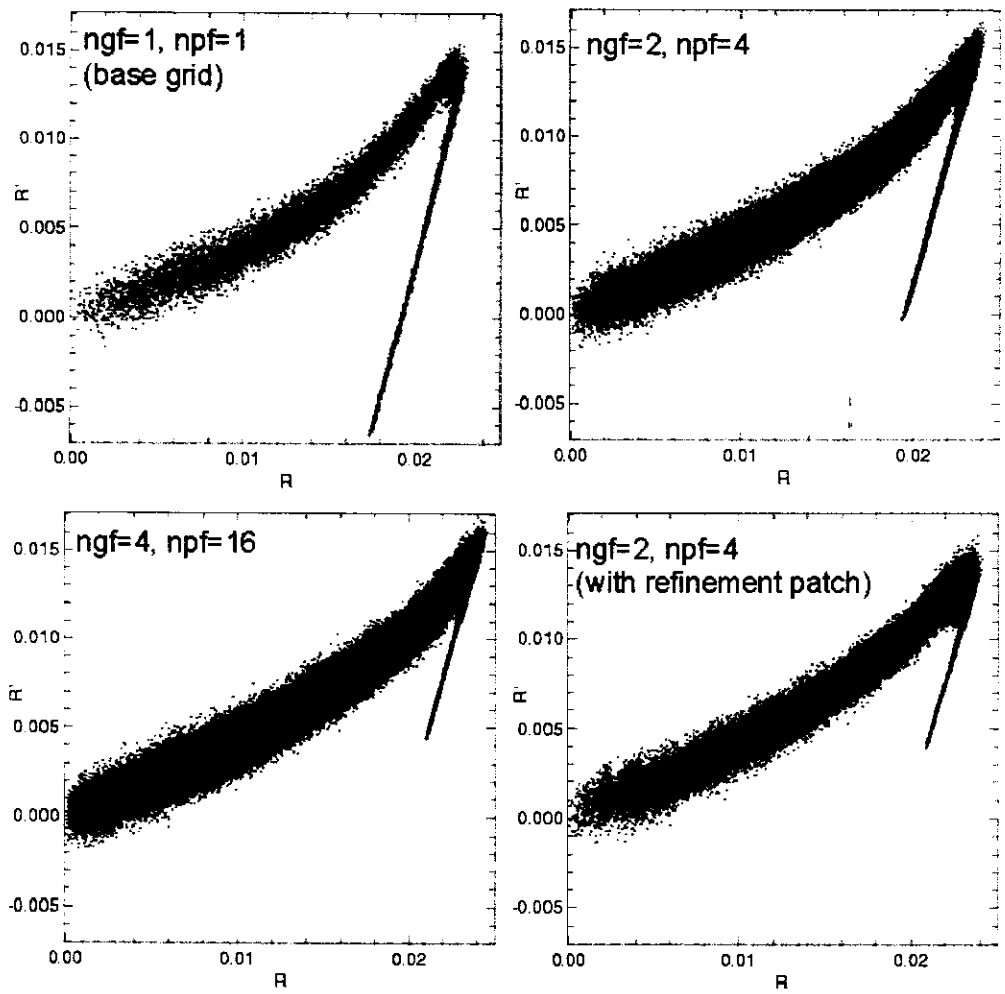

Fig. 5. Phase-space projections of the beam slice distribution in the r-r' plane at the end of the runs, for 39 cm $< z <$ 40 cm (see text for runs labeling description). A higher resolution leads to a smaller hook. Applying a mesh refinement around the emitting area allows us to recover the correct length of the hook at a reduced cost (factor of 4 in this case).

executable routines is provided to the WARP linker which merges the two packages together. Appropriate calls to Chombo routines are made by WARP FORTRAN's routines which are triggered by a flag which is set up by the user in WARP's Python script interface. For specialized use, some of the Chombo routines, such as its AMR Poisson solver, are callable directly from WARP's Python interface.

While the production-level general three-dimensional AMR Poisson solver is being developed in the Chombo package, we have built a prototype axisymmetric (RZ) AMR–Poisson solver on the foundations of WARP's r-z Poisson solver. This has allowed us to begin to explore the benefits of AMR for injector simulation, as described in the next section.

5. EXAMPLE: SIMULATION OF THE HCX SOURCE WITH WARP-RZ

Figure 2 shows a snapshot taken from a movie of an end-to-end HCX simulation. As explained in Haber *et al.* (2002), a successful end-to-end simulation requires a very detailed simulation of the front end. We present some results obtained on the convergence study of the HCX source simulation. To focus on the source, no quadrupole were considered, which allowed us to benefit from the axisymmetry of the system, thus reducing the dimensionality of the simulations. Assuming that we are interested only in the steady-state solution, the simulations were performed in a "quasi time-dependent" mode, where the electrostatic field was solved every 10 time steps. In all the runs, we have used a time step of 1 ns and the run was stopped after 1000 time steps (1 μs physical time). A snapshot of the particle distribution at the end of the run is given in Figure 3 for a calculation using a grid of $nr \times nz = 56 \times 640$ cells (giving $\delta x \approx 3.6$ mm and $\delta z \approx 0.6$ mm). The number of cells was chosen so that there are at least 10 grid lines crossing the emitting region in each dimension (the emitting surface has a radius of 5.08 cm and an extension of 6.45 mm in z). We will refer to this run as the "base run." We define the simulation parameters of other runs by two integers, *ngf* and *npf*, which represent, respec-

tively, the multiplicative factor for the number of cells in each dimension and the number of particles, with respect to the base run. Hence, a run labeled $\{ngf = 2, npf = 4\}$ had a grid of 112×1280 cells and used four times the number of macroparticles of the base run. We ran the cases $\{ngf = 1, npf = 1\}$ (i.e., base grid), $\{ngf = 2, npf = 4\}$, $\{ngf = 4, npf = 16\}$, and $\{ngf = 2, npf = 4$, with mesh refinement$\}$. For the last case, a mesh refinement patch having two times the resolution of the initial grid in each dimension was applied on the emitter region. The limits of the patch are shown by a red line on Figure 3. Let cc be the computing cost of a run; then we have: $cc(ngf = 4, npf = 16) = 4cc(ngf = 2, npf = 4) = 16cc(ngf = 1, npf = 1) \approx 4cc(ngf = 2, npf = 4$, with mesh refinement). Figure 4 and Figure 5 display, respectively, the emittance as a function of z and a beam slice phase-space projection, taken at the end of the runs. From Figure 4, it follows that higher resolution means a reduction of the emittance. From Figure 5, we can infer that this emittance reduction is linked to the length of a hook at the edge of the distribution, which also reduces when the resolution rises. We remark that we obtained results very close to the highest resolution case $\{ngf = 4, npf = 16\}$ with the case $\{ngf = 2, npf = 4$, with mesh refinement$\}$, at a fourth of the computational cost.

6. CONCLUSION

We have discussed the potential benefits of the introduction of the AMR technique for the heavy ion fusion program, as well as the difficulties that arise in its application to plasma and accelerator modeling. We have presented the ongoing effort of coupling an existing AMR package with the PIC code WARP and have shown that a significant reduction factor could be obtained in computer time cost on the simulation of a key issue for HIF. We conclude from this preliminary study that the introduction of the AMR technique into beam and plasma simulations offers the potential of more efficient calculations, leading us to reach the goal of integrated end-to-end simulation significantly sooner than would be possible otherwise.

REFERENCES

Bell, J.B., Berger, J.S., Saltzman, J.S. & Welcome, M. (1991). Three dimensional adaptive mesh refinement for hyperbolic conservation laws. LLNL Report UCRL-JC-108794. Livermore, CA: Lawrence Livermore National Laboratory.

Berenger, J.-P. (1994). A perfectly matched layer for the absorption of electromagnetic waves. *J. Comput. Phys.* **114**, 185.

Berger, M.J. & Oliger, J. (1984). Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.* **53**, 484.

Berger, M.J. & Colella, P. (1989). Local adaptive mesh refinement for shock hydrodymanics. *J. Comput. Phys.* **82**, 64.

Collino, F., Fouquet, T. & Joly, P. (1999). Une methode de raffinement de maillage espace-temps pour le systeme de Maxwell en dimension un. *C.R. Acad. Sci. Paris* **328**, 263.

Engquist, B. & Majda, A. (1977). Absorbing boundary conditions for the numerical simulation of waves. *Math. Comput.* **31**, 629.

Friedman, A. (1990). A 2nd-order implicit particle mover with adjustable damping. *J. Comput. Phys.* **90**, 292.

Grote, D.P., Friedman, A. & Haber, I. (1996). Three-dimensional simulations of high current beams in induction accelerators with WARP3d. *Fusion Eng. Des.* **32–33**, 193.

Haber, I., Bieniosek, F.M., Celata, C.M., Friedman, A., Grote, D.P., Henestroza, E., Vay, J.-L., Bernal, S., Kishek, R.A., O'Shea, P.G., Reiser, M. & Herrmannsfeldt, W.B. (2002). End-to-end simulation: The front end. *Laser Part. Beams* **20**, 431–433.

Vay, J.-L. (2001). An extended FDTD scheme for the wave equation: Application to multiscale electromagnetic simulation. *J. Comput. Phys.* **167**, 72.

Vay, J.-L. (2002). Asymmetric perfectly matched layer for the absorption of waves. *J. Comput. Phys.* **183**, 367.