SOME COMPUTATIONAL CHALLENGES IN THE MODELING OF ACCELERATORS AND THEIR SOLUTIONS IN THE SIMULATION CODE WARP

J.-L. Vay*, C. Benedetti, LBNL, USA
D.P. Grote, R.H. Cohen, A. Friedman LLNL, USA

Abstract

This paper presents an overview of the Particle-In-Cell accelerator code Warp's capabilities, summarizing recent original numerical methods that were developed within Warp, including a large-timestep "drift-Lorentz" mover for arbitrarily magnetized species, a relativistic Lorentz invariant leapfrog particle pusher, an electromagnetic solver with tunable numerical dispersion and efficient stride-based digital filtering, Particle-In-Cell with Adaptive Mesh Refinement, and simulations in Lorentz boosted frames.

INTRODUCTION

The Particle-In-Cell (PIC) Framework Warp [1] was originally developed to simulate space-charge-dominated beam dynamics in induction accelerators for heavy-ion fusion (HIF) [2]. In recent years, the physics models in the code have been generalized, so that Warp can model beam injection, complicated boundary conditions, denser plasmas, a wide variety of accelerator lattice components, and the non-ideal physics of beams interacting with walls and plasmas. The code now has an international user base and is being applied to projects both within and far removed from the HIF community. Ongoing or recent examples of applications outside HIF include the modeling of plasma traps for the production of anti-Hydrogen [3], Paul traps [4, 5], non-conventional Penning-Malmberg micro-trap [6], transport of electron beams in the UMER ring [7], ECR ion sources [8], capture and control of laser-accelerated proton beams [9], and fundamental studies of multipacting [10]. It is also applied to the study and design of existing and next generation high-energy accelerators including the study of electron cloud effects [11], coherent synchrotron radiation [12] and laser wakefield acceleration [13].

These studies have necessitated the introduction or development of advanced numerical methods, including methods to model multiple-species effects in accelerators and chambers, efficient ensemble methods, particle advance algorithms that allow a longer time step, and adaptive mesh refinement (AMR).

NOVEL ALGORITHMS

Hybrid Drift-Lorentz

It was observed in [14] that the Boris pusher causes particles to gyrate with spuriously large radius for time steps that are large compared to the gyroperiod, albeit with the correct drift velocities (provided the gradients are still sampled adequately). A new solver that interpolates between the Boris velocity push and a drift kinetic advance was developed and implemented in Warp [15, 17]; it reproduces both the correct drift velocity and gyroradius for an arbitrarily large ratio of time step δt relative to cyclotron period τ_c, as well as correct detailed orbit dynamics in the small-timestep limit. The pusher has provided an order of magnitude or more saving in computing resources in the simulations of electron cloud effects in the HCX experiment [15, 16]. An implicit time-advance scheme incorporating drift-Lorentz interpolation has also been developed [17].

Lorentz Invariant Advance

The relativistic version of the Boris (or Hybrid Lorentz-Drift) particle pusher does not maintain strict Lorentz invariance, resulting eventually in unacceptably large inaccuracies when modeling the transport of ultra-relativistic beams in accelerators. To this effect, an alternative to the Boris pusher that conserves strict Lorentz invariance (to machine precision) was developed and implemented in Warp, and its effectiveness demonstrated on the modeling from first principles of the interaction of a 500 GeV proton beam with a background of electrons [18]. The pusher has subsequently been implemented by others and has also proven useful for correctly capturing the drift speed of electrons of a highly magnetized relativistic electron-ion flow in astrophysical simulations using the code TRISTAN [19].

Electromagnetic Solver

Warp's electromagnetic solver is based on the Non-Standard Finite-Difference (NSFD) technique [20, 21], which is an extension of the Finite-Difference Time-Domain technique to larger stencils in the plane perpendicular to the direction of the finite difference. This gives the user some control on the numerical dispersion and Courant time step limits which do depend on those parameters. As shown in [22], for a given set of parameters, and for cubic cells, the Courant time step multiplied by the speed of light equals the cell size, and the numerical dispersion vanishes along the main axes. More details on the solver implementation and characteristics for several sets of coefficients are available in [23]. Also described in [23] are the implementation of Perfectly Matched Layers for the absorption of waves at grid boundaries and of Friedman's damping algorithm for noise control [24]. In the same paper, it is shown that introducing a stride in the usage of standard linear filtering allows for construction of efficient iterative sideband

* jlvay@lbl.gov

ISBN 978-3-95450-116-8

digital filters that are nonetheless compact, thus well suited for implementation on parallel computers.

Mesh Refinement

The mesh refinement methods that have been implemented in Warp were developed following the following principles: i) avoidance of spurious effects from mesh refinement, or minimization of such effects; ii) user controllability of the spurious effects' relative magnitude; iii) simplicity of implementation. The two main generic issues that were identified are: a) spurious self-force on macroparticles close to the mesh refinement interface [25, 26]; b) reflection (and possible amplification) of short wavelength electromagnetic waves at the mesh refinement interface [27]. The two effects are due to the loss of translation invariance introduced by the asymmetry of the grid on each side of the mesh refinement interface.

A cornerstone of the Particle-In-Cell method is that assuming a particle lying in a hypothetical infinite grid, then if the grid is regular and symmetrical, and if the order of field gathering matches the order of charge (or current) deposition, then there is no self-force of the particle acting on itself: a) anywhere if using the so-called "momentum conserving" gathering scheme; b) on average within one cell if using the "energy conserving" gathering scheme [28]. A breaking of the regularity and/or symmetry in the grid, whether it is from the use of irregular meshes or mesh refinement, and whether one uses finite difference, finite volume or finite elements, results in a net spurious self-force (which does not average to zero over one cell) for a macroparticle close to the point of irregularity (mesh refinement interface for the current purpose) [25, 26].

A method was devised and implemented in Warp for reducing the magnitude of spurious self-forces near the coarse-fine boundaries as follows. Noting that the coarse grid solution is unaffected by the presence of the patch and is thus free of self-force, extra "transition" cells are added around the "effective" refined area. Within the effective area, the particles gather the potential in the fine grid. In the extra transition cells surrounding the refinement patch, the force is gathered directly from the coarse grid. The number of cells allocated in the transition zones is controllable by the user in Warp, giving the opportunity to check whether the spurious self-force is affecting the calculation by repeating it using different thicknesses of the transition zones. Automatic remeshing has been implemented in Warp following the procedure described in [29], refining on criteria based on measures of local charge density magnitude and gradients.

The method that is used for electrostatic mesh refinement is not directly applicable to electromagnetic calculations. As was shown in section 3.4 of [30], refinement schemes relying solely on interpolation between coarse and fine patches lead to the reflection with amplification of the short wavelength modes that fall below the cutoff of the Nyquist frequency of the coarse grid. Unless these modes are damped heavily or prevented from occurring at their source, they may affect particle motion and their effect can escalate if trapped within a patch, via multiple successive reflections with amplification.

To circumvent this issue, an additional coarse patch (with the same resolution as the parent grid) is added, as described in [31]. Both the fine and the coarse grid patches are terminated by Perfectly Matched Layers, reducing wave reflection by orders of magnitude, controllable by the user [32, 33].

The source current resulting from the motion of charged macroparticles within the refined region is accumulated on the fine patch and is then interpolated onto the coarse patch and added onto the parent grid. The process is repeated recursively from the finest level down to the coarsest. The Maxwell equations are then solved for one time interval on the entire set of grids, by default for one time step using the time step of the finest grid. The field on the coarse and fine patches only contain the contributions from the particles that have evolved within the refined area but not from the current sources outside the area. The total contribution of the field from sources within and outside the refined area is obtained by adding the field from the refined grid $F(r)$, and adding an interpolation I of the difference between the relevant subset s of the field in the parent grid $F(s)$ and the field of the coarse grid $F(c)$, on an auxiliary grid a, i.e. $F(a) = F(r) + I[F(s) - F(c)]$. The field on the parent grid subset $F(s)$ contains contributions from sources from both within and outside of the refined area. Thus, in effect, there is substitution of the coarse field resulting from sources within the patch area by its fine resolution counterpart. The operation is carried out recursively starting at the coarsest level up to the finest.

Lorentz Boosted Frame

A method was recently proposed to speed up full PIC simulations of a certain class of relativistic interactions by performing the calculation in a Lorentz boosted frame [34], taking advantage of the properties of space/time contraction and dilation of special relativity to render space and time scales (that are separated by orders of magnitude in the laboratory frame) commensurate in a Lorentz boosted frame, resulting in far fewer computer operations. The method has been applied successfully to the modeling of laser plasma acceleration [35, 36, 37, 13], electron cloud effects [18], free electron lasers [38], coherent synchrotron radiation [12], and production of ultrabright attosecond x-ray pulses [39].

In a laser plasma accelerator, a laser pulse is injected through a plasma, creating a wake of regions with very strong electric fields of alternating polarity [40]. An electron beam that is injected with the appropriate phase can thus be accelerated to high energy in a distance that is much shorter than with conventional acceleration techniques [41]. The simulation of a laser plasma acceleration stage from first principles using the Particle-In-Cell

ISBN 978-3-95450-116-8

technique in the laboratory frame is very demanding computationally, as the evolution of micron-scale long laser oscillations needs to be followed over millions of time steps as the laser pulse propagates through a meter long plasma for a 10 GeV stage. In the laboratory frame the laser pulse is much shorter than the wake, whose wavelength is also much shorter than the acceleration distance ($\lambda_{laser} \ll \lambda_{wake} \ll \lambda_{acceleration}$). In a Lorentz boosted frame moving at a speed near the speed of light with the laser in the plasma, the laser will be Lorentz expanded (by a factor $(1 + v_f/c)\gamma_f$ where $\gamma_f = (1 - v_f^2/c^2)^{-1/2}$ and v_f is the velocity of the frame and c is the speed of light). The plasma (now moving opposite to the incoming laser at velocity $-v_f$) is Lorentz contracted (by a factor γ_f). In a boosted frame moving with the wake ($\gamma_f \approx \gamma_{wake}$), the laser wavelength, the wake and the acceleration length are now commensurate ($\lambda_{laser} < \lambda_{wake} \approx \lambda_{acceleration}$), leading to far fewer time steps by a factor $(1 + v_f/c)^2\gamma_f^2$, hence computer operations [34, 13].

A speedup of up to a million times was reported for Warp modeling of a hypothetical 1 TeV stage [42]. Control of a violent numerical instability (which nature is being investigated) that had been plaguing early attempts was obtained via the combination of: (i) the use of Warp's tunable electromagnetic solver and efficient wideband filtering [23], (ii) observation of the benefits of hyperbolic rotation of spacetime on the laser spectrum in boosted frame simulations [42], and (iii) identification of a special time step at which the growth rate of the instability is greatly reduced [23]. In addition, a novel numerical method for injecting the laser pulse through a moving planar antenna was introduced in Warp [13].

CONCLUSION

The Warp code-framework has recently been augmented with various novel methods including PIC with adaptive mesh refinement, a large-timestep mover for particles of arbitrary magnetized species, a new relativistic Lorentz invariant leapfrog particle pusher, simulations in Lorentz boosted frames, an electromagnetic solver with tunable numerical dispersion and efficient stride-based digital filtering. With its new capabilities and thanks to a design that allows for a high degree of versatility, the range of application of Warp has considerably widened far beyond the initial application to the Heavy Ion Fusion Science program.

ACKNOWLEDGMENTS

Work supported by US-DOE Contracts DE-AC02-05CH11231 and DE-AC52-07NA27344, the US-DOE SciDAC program ComPASS and the US-LHC Accelerator Research Program (LARP). Used resources of NERSC, supported by US-DOE Contract DE-AC02-05CH11231.

REFERENCES

[1] D. Grote, *et al.*, AIP Conf. Proc., no. 749, 2005, pp. 55–8.
[2] B. Logan, *et al.*, Nucl. Instr. & Meth. A 577 (2007) 1–7.
[3] K. Gomberoff, *et al.*, Phys. Plasmas 14 (10).
[4] E. P. Gilson, *et al.*, Phys. Plasmas 17 (5)
[5] S. Ohtsubo, *et al.* Phys. Rev. ST-AB 13 (4).
[6] P. Folegati, *et al.*, J. of Physics: Conf. Series 262 (2011).
[7] I. Haber, *et al.* Nucl. Instr. & Methods A 606 (2009) 64–68
[8] D. Winklehner, *et al.* J. Instr. 5 (2010).
[9] F. Nürnberg *et al.*, J. Physics: Conf. Series 244 (2010).
[10] R. A. Kishek, *et al.*, Phys. Rev. Lett. 108 (2012) 035003.
[11] J.-L. Vay, *et al.*, Proc. Part. Accel. Conf. (2011).
[12] W. Fawley, *et al.* Proc. Int. Part. Accel. Conf., (2010).
[13] J. L. Vay, *et al.*, Phys. Plasmas 18 (12).
[14] S. Parker, *et al.*, 97 (1) (1991) 91–102.
[15] R. H. Cohen, *et al.*, Phys. Plasmas 12 (5) (2004).
[16] A. W. Molvik, *et al.*, Phys. Plasmas 14 (2006).
[17] R. H. Cohen, *et al.*, Nucl. Instr. & Meth. A 606 (2009) 53.
[18] J. L. Vay, Phys. Plasmas 15 (2008) 056701.
[19] L. Sironi, A. Spitkovsky, private Communication (2011).
[20] J. Cole, IEEE Trans. Microwave Th. & Tech. 45 (1997) 991.
[21] J. Cole, IEEE Trans. Ant. & Prop. 50 (2002) 1185–1191.
[22] M. Karkkainen, *et al.*, ICAP, Chamonix, France, (2006).
[23] J. L. Vay, *et al.*, J. Comput. Phys. 230 (2011) 5908–5929.
[24] A. Friedman, J. Comput. Phys. 90 (1990) 292–312.
[25] J. Vay, *et al.*, Laser & Part. Beams 20 (2002) 569–575.
[26] P. Colella, P. C. Norgaard, J. Comput. Phys. 229 (2010) 947.
[27] J.-L. Vay, J. Comput. Phys. 165 (2) (2000) 511–521.
[28] C. Birdsall, A. Langdon, Plasma physics via computer simulation, Adam-Hilger, 1991.
[29] J.-L. Vay, *et al.* Nucl. Instr. & Meth. A 544 (2005) 347–352
[30] J.-L. Vay, J. Comput. Phys. 167 (2001) 72–98.
[31] J.-L. Vay, *et al.*, Comput. Phys. Comm. 164 (2004) 171–177.
[32] J. Berenger, J. Comput. Phys. 127 (1996) 363–379.
[33] J.-L. Vay, J. Comput. Phys. 183 (2002) 367–399.
[34] J.-L. Vay, Phys. Rev. Lett. 98 (2007) 130405.
[35] D. Bruhwiler, *et al.* AIP Conf. Proc., 1086, (2009), 29–37.
[36] J.-L. Vay, *et al.*, Proc. PAC09, Vancouver, Canada (2009).
[37] S. F. Martins, *et al.*, Nat. Physics 6 (2010) 311–316.
[38] W. Fawley, J.-L. Vay, Proc. FEL 2010, Sweden (2010).
[39] H. C. Wu, *et al.* Phys. Rev. ST-AB 14 (2011).
[40] T. Tajime, J. Dawson, Phys. Rev. Let. 43 (1979) 267.
[41] W. P. Leemans, *et al.* Nat. Physics 2 (2006) 696–699.
[42] J. Vay, *et al.*, Phys. Plasmas 18 (2011) 030701.

ISBN 978-3-95450-116-8