

HIGH-PERFORMANCE MODELING OF PLASMA-BASED ACCELERATION AND LASER-PLASMA INTERACTIONS*

J.-L. Vay[†], G. Blaclard, R. Lehe, M. Lobet, LBNL, Berkeley, CA, U.S.A.
 M. Kirchen, University of Hamburg, Hamburg, Germany
 H. Vincenti, Commissariat a l'Energie Atomique, Gif-Sur-Yvette, France
 P. Lee, Université Paris-Saclay, Orsay, France
 B. B. Godfrey, U. Maryland, College Park, MD, USA

Abstract

Large-scale numerical simulations are essential to the design of plasma-based accelerators and laser-plasma interactions for ultra-high intensity (UHI) physics. The electromagnetic Particle-In-Cell (PIC) approach is the method of choice for self-consistent simulations, as it is based on first principles, and captures all kinetic effects, and also scales easily (for uniform plasmas) to many cores on super-computers. The standard PIC algorithm relies on second-order finite-difference discretizations of the Maxwell and Newton-Lorentz equations. We present here novel PIC formulations, based on the use of analytical pseudo-spectral Maxwell solvers, which enable near-total elimination of the numerical Cherenkov instability and increased accuracy over the standard PIC method. We also discuss the latest implementations in the PIC modules Warp-PICSAR and FBPIC on the Intel Xeon Phi and GPU architectures. Examples of applications are summarized on the simulation of high-harmonic generation with plasma mirrors and of laser-plasma accelerators.

PSEUDO SPECTRAL ANALYTICAL TIME DOMAIN (PSATD)

Maxwell's equations in Fourier space are given by

$$\frac{\partial \tilde{\mathbf{E}}}{\partial t} = i\mathbf{k} \times \tilde{\mathbf{B}} - \tilde{\mathbf{J}} \quad (1a)$$

$$\frac{\partial \tilde{\mathbf{B}}}{\partial t} = -i\mathbf{k} \times \tilde{\mathbf{E}} \quad (1b)$$

$$[i\mathbf{k} \cdot \tilde{\mathbf{E}} = \tilde{\rho}] \quad (1c)$$

$$[i\mathbf{k} \cdot \tilde{\mathbf{B}} = 0] \quad (1d)$$

where \tilde{a} is the Fourier Transform of the quantity a . As with the real space formulation, provided that the continuity equation $\partial \tilde{\rho} / \partial t + i\mathbf{k} \cdot \tilde{\mathbf{J}} = 0$ is satisfied, then the last two equations will automatically be satisfied at any time if satisfied initially and do not need to be explicitly integrated.

Decomposing the electric field and current between longitudinal and transverse components $\tilde{\mathbf{E}} = \tilde{\mathbf{E}}_L + \tilde{\mathbf{E}}_T =$

$\hat{\mathbf{k}}(\hat{\mathbf{k}} \cdot \tilde{\mathbf{E}}) - \hat{\mathbf{k}} \times (\hat{\mathbf{k}} \times \tilde{\mathbf{E}})$ and $\tilde{\mathbf{J}} = \tilde{\mathbf{J}}_L + \tilde{\mathbf{J}}_T = \hat{\mathbf{k}}(\hat{\mathbf{k}} \cdot \tilde{\mathbf{J}}) - \hat{\mathbf{k}} \times (\hat{\mathbf{k}} \times \tilde{\mathbf{J}})$ gives

$$\frac{\partial \tilde{\mathbf{E}}_T}{\partial t} = i\mathbf{k} \times \tilde{\mathbf{B}} - \tilde{\mathbf{J}}_T \quad (2a)$$

$$\frac{\partial \tilde{\mathbf{E}}_L}{\partial t} = -\tilde{\mathbf{J}}_L \quad (2b)$$

$$\frac{\partial \tilde{\mathbf{B}}}{\partial t} = -i\mathbf{k} \times \tilde{\mathbf{E}} \quad (2c)$$

with $\hat{\mathbf{k}} = \mathbf{k}/k$.

If the sources are assumed to be constant over a time interval Δt , the system of equations is solvable analytically and is given by (see [1] for the original formulation and [2] for a more detailed derivation):

$$\tilde{\mathbf{E}}^{n+1} = C\tilde{\mathbf{E}}^n + iS\hat{\mathbf{k}} \times \tilde{\mathbf{B}}^n - \frac{S}{k}\tilde{\mathbf{J}}^{n+1/2} \quad (3a)$$

$$+ (1-C)\hat{\mathbf{k}}(\hat{\mathbf{k}} \cdot \tilde{\mathbf{E}}^n) + \hat{\mathbf{k}}(\hat{\mathbf{k}} \cdot \tilde{\mathbf{J}}^{n+1/2}) \left(\frac{S}{k} - \Delta t \right), \quad (3b)$$

$$\tilde{\mathbf{B}}^{n+1} = C\tilde{\mathbf{B}}^n - iS\hat{\mathbf{k}} \times \tilde{\mathbf{E}}^n \quad (3c)$$

$$+ i\frac{1-C}{k}\hat{\mathbf{k}} \times \tilde{\mathbf{J}}^{n+1/2}. \quad (3d)$$

with $C = \cos(k\Delta t)$ and $S = \sin(k\Delta t)$.

For fields generated by the source terms without the self-consistent dynamics of the charged particles, this algorithm is free of numerical dispersion and is not subject to a Courant condition. Furthermore, this solution is exact for any time step size subject to the assumption that the current source is constant over that time step.

ALTERNATE FORMULATION IN A GALILEAN FRAME

The numerical Cherenkov instability (NCI) [3] is the most serious numerical instability affecting multidimensional PIC simulations of relativistic particle beams and streaming plasmas [4–9]. It arises from coupling between possibly numerically distorted electromagnetic modes and spurious beam modes, the latter due to the mismatch between the Lagrangian treatment of particles and the Eulerian treatment of fields [10].

A new scheme was recently proposed, in [11, 12], which completely eliminates the NCI for a plasma drifting at a

* Work supported by US-DOE Contracts DE-AC02-05CH11231 and by the European Commission through the Marie Slowdoska-Curie actions (Marie Curie IOF fellowship PICSAR grant number 624543). Used resources of NERSC, supported by US-DOE Contract DE-AC02-05CH11231.

[†] jlway@lbl.gov

uniform relativistic velocity – with no arbitrary correction – by simply integrating the PIC equations in *Galilean coordinates* (also known as *comoving coordinates*). More precisely, in the new method, the Maxwell equations in *Galilean coordinates* are integrated analytically, using only natural hypotheses, within the PSATD framework.

The idea of the proposed scheme is to perform a Galilean change of coordinates, and to carry out the simulation in the new coordinates $\mathbf{x}' = \mathbf{x} - \mathbf{v}_{gal}t$ where $\mathbf{x} = x\mathbf{u}_x + y\mathbf{u}_y + z\mathbf{u}_z$ and $\mathbf{x}' = x'\mathbf{u}_x + y'\mathbf{u}_y + z'\mathbf{u}_z$ are the position vectors in the standard and Galilean coordinates respectively.

When choosing $\mathbf{v}_{gal} = \mathbf{v}_0$, where \mathbf{v}_0 is the speed of the bulk of the relativistic plasma, the plasma does not move with respect to the grid in the Galilean coordinates \mathbf{x}' – or, equivalently, in the standard coordinates \mathbf{x} , the grid moves along with the plasma. The heuristic intuition behind this scheme is that these coordinates should prevent the discrepancy between the Lagrangian and Eulerian point of view, which gives rise to the NCI [10].

An important remark is that the Galilean change of coordinates is a simple translation. Thus, when used in the context of Lorentz-boosted simulations [13], it does readily preserve the relativistic dilatation of space and time which gives rise to the characteristic computational speedup of the boosted-frame technique.

In the Galilean coordinates \mathbf{x}' , the equations of particle motion and the Maxwell equations take the form

$$\frac{d\mathbf{x}'}{dt} = \frac{\mathbf{p}}{\gamma m} - \mathbf{v}_{gal} \quad (4a)$$

$$\frac{d\mathbf{p}}{dt} = q \left(\mathbf{E} + \frac{\mathbf{p}}{\gamma m} \times \mathbf{B} \right) \quad (4b)$$

$$\left(\frac{\partial}{\partial t} - \mathbf{v}_{gal} \cdot \nabla' \right) \mathbf{B} = -\nabla' \times \mathbf{E} \quad (4c)$$

$$\frac{1}{c^2} \left(\frac{\partial}{\partial t} - \mathbf{v}_{gal} \cdot \nabla' \right) \mathbf{E} = \nabla' \times \mathbf{B} - \mu_0 \mathbf{J} \quad (4d)$$

where ∇' denotes a spatial derivative with respect to the Galilean coordinates \mathbf{x}' .

Integrating these equations from $t = n\Delta t$ to $t = (n+1)\Delta t$ results in the following update equations (see [12] for the details of the derivation):

$$\begin{aligned} \tilde{\mathbf{B}}^{n+1} &= \theta^2 C \tilde{\mathbf{B}}^n - \frac{\theta^2 S}{ck} i\mathbf{k} \times \tilde{\mathbf{E}}^n \\ &+ \frac{\theta \chi_1}{\epsilon_0 c^2 k^2} i\mathbf{k} \times \tilde{\mathbf{J}}^{n+1/2} \end{aligned} \quad (5a)$$

$$\begin{aligned} \tilde{\mathbf{E}}^{n+1} &= \theta^2 C \tilde{\mathbf{E}}^n + \frac{\theta^2 S}{k} ci\mathbf{k} \times \tilde{\mathbf{B}}^n \\ &+ \frac{i\nu\theta\chi_1 - \theta^2 S}{\epsilon_0 ck} \tilde{\mathbf{J}}^{n+1/2} \\ &- \frac{1}{\epsilon_0 k^2} \left(\chi_2 \hat{\rho}^{n+1} - \theta^2 \chi_3 \hat{\rho}^n \right) i\mathbf{k} \end{aligned} \quad (5b)$$

where we used the short-hand notations $\tilde{\mathbf{E}}^n \equiv \tilde{\mathbf{E}}(\mathbf{k}, n\Delta t)$, $\tilde{\mathbf{B}}^n \equiv \tilde{\mathbf{B}}(\mathbf{k}, n\Delta t)$ as well as:

$$C = \cos(ck\Delta t) \quad S = \sin(ck\Delta t) \quad k = |\mathbf{k}| \quad (6a)$$

$$\nu = \frac{\mathbf{k} \cdot \mathbf{v}_{gal}}{ck} \quad \theta = e^{i\mathbf{k} \cdot \mathbf{v}_{gal}\Delta t/2} \quad \theta^* = e^{-i\mathbf{k} \cdot \mathbf{v}_{gal}\Delta t/2} \quad (6b)$$

$$\chi_1 = \frac{1}{1 - \nu^2} (\theta^* - C\theta + i\nu\theta S) \quad (6c)$$

$$\chi_2 = \frac{\chi_1 - \theta(1 - C)}{\theta^* - \theta} \quad \chi_3 = \frac{\chi_1 - \theta^*(1 - C)}{\theta^* - \theta} \quad (6d)$$

As shown in [11, 12], the elimination of the NCI with the new Galilean integration is verified empirically via PIC simulations of uniform drifting plasmas and laser-driven plasma acceleration stages, and confirmed by a theoretical analysis of the instability.

EXTENSION TO QUASI-CYLINDRICAL GEOMETRY

When the geometry of the simulated system is cylindrically-symmetric, it is possible to take advantage of the symmetry of the problem to reduce the computational cost of the algorithm [14–17]. Without loss of generality, fields \mathbf{E} , \mathbf{B} , \mathbf{J} and ρ in cylindrical coordinates (r, θ, z) , can be expressed as Fourier series in θ :

$$F(r, \theta, z) = \text{Re} \left[\sum_{\ell=0}^{\infty} \tilde{F}_{\ell}(r, z) e^{-i\ell\theta} \right] \quad (7)$$

$$\text{with} \quad \tilde{F}_{\ell} = C_{\ell} \int_0^{2\pi} d\theta F(r, \theta, z) e^{i\ell\theta} \quad (8)$$

$$\text{and} \quad \begin{cases} C_0 = 1/2\pi \\ C_{\ell} = 1/\pi \quad \text{for } \ell > 0 \end{cases} \quad (9)$$

where F represents any of the field quantities, and where the \tilde{F}_{ℓ} are the associated Fourier components (ℓ is the index of the corresponding azimuthal mode). In the case of a cylindrically-symmetric laser pulse, only the very first modes have non-zero components. For instance, the wakefield of a laser plasma accelerator is represented exclusively by the mode $\ell = 0$. (This is because the field quantities associated with the wakefield are independent of θ .) On the other hand, the field of the laser pulse *does* depend on θ , in cylindrical coordinates.

The PSATD solver was ported to this geometry in the code FBPIC (Fourier-Bessel Particle- In-Cell), enabling a quasi-cylindrical solver with no numerical dispersion not Courant condition [18]. In addition, as explained and demonstrated in [11, 12], the Galilean PSATD presented in the preceding section applies readily to this geometry and suppresses the numerical Cherenkov instability as effectively as in the Cartesian case.

EFFICIENT IMPLEMENTATION ON MANICORE CPU AND GPU

The set of electromagnetic PIC kernel subroutines of the PIC code Warp [19] were turned into a standalone software

package named PICSAR (for Particle-In-Cell Scalable Architecture Resources). The subroutines were optimized for the latest Intel Xeon Phi architecture, with a novel portable vectorization algorithm, and parallelization intra- and inter-node based state-of-the-art OpenMP directives and MPI communications [20]. A new class was implemented that enables a seamless use of the optimized PICSAR subroutines in Warp, with minimal modifications to an existing Warp input script (only a couple of lines need to be modified). In addition, PICSAR brings efficient dynamic load balancing to electromagnetic PIC simulations. Warp and PICSAR have been ported to the U.S. Department of Energy supercomputers Edison, Cori and Mira.

The code FBPIC is written in Python [18]. For performance, this implementation makes use of the pre-compiled libraries FFTW [21] and BLAS [22] (for the matrix multiplication in the DHT), and utilizes the Numba just-in-time compiler [23] for the computationally-intensive parts of the code (current deposition and field gathering). In addition, the code was developed for both single-CPU and single-GPU architectures (with the GPU runs being typically more than 40 times faster than the equivalent CPU runs, on modern hardware, as compared to a run using only one CPU). Importantly, a precise timing of the different routines showed that, although the time taken by the spectral transforms (FFT and DHT) is not entirely negligible, it does not usually dominate the PIC cycle. For instance, on a K20 GPU and for a grid with $N_z = 4096$, $N_r = 256$ and 16 particles per cell, the FFTs and DHTs take up 6% and 14% of the PIC cycle respectively, while the rest of the time is dominated by the field gathering and current deposition.

APPLICATION OF PSATD TO THE STUDY OF HIGH HARMONIC GENERATION AND LASER PLASMA ACCELERATION

Standard FDTD Maxwell solvers fail to accurately describe Doppler harmonic generation with Particle-In-Cell codes, because they introduce numerical dispersion, causing angular deviation that significantly affects the harmonic spectrum. In [24], we showed that this angular deviation can be understood as a simple refraction of harmonic beams when these enter vacuum at the plasma-vacuum interface. A simple model based on Snell-Descartes law was derived that can accurately reproduce the angular deviation observed in PIC simulations. This model can now be used to estimate the minimum resolution required to avoid this spurious deviation. The results of our model show that the required computing resources in 3D for these standard solvers exceed by far the ones of current petascale and future exascale supercomputer capacities. In that case, the solution is to use dispersion-free pseudo-spectral solvers (as PSATD), for which there is no angular deviation even at moderate resolution.

In the past, we have carried very detailed studies of the convergence of Lorentz boosted frame PIC simulations of laser plasma accelerator stages in the linear regime ($a_0=1$)

with external injection, demonstrating convergence to the percent level or better. While some agreement has been reported by other groups in the modeling of LPA in the non-linear regime ($a_0>2$), a good quantitative agreement has not yet been reported. We have thus carried out a campaign of simulations to study the convergence of the properties of the electron bunch generated via self-injection scheme using Lorentz boosted frames in the non-linear regime. Several sets of physical and numerical parameters, as well as numerical schemes, are being considered, including FDTD and PSATD in Cartesian and quasi-cylindrical coordinates. Preliminary results demonstrate good convergence of the beam moments. Further studies are underway to demonstrate convergence of higher energy stages using lower plasma densities.

CONCLUSION

The development of very accurate pseudo-spectral electromagnetic solvers in Cartesian and quasi-cylindrical geometries enables simulations that are free of numerical dispersion in vacuum. In addition, the PSATD formulation has been extended to an implementation of PSATD-PIC in a Galilean frame that is free of numerical Cherenkov for simulations of plasmas drifting at a uniform relativistic velocities, such as for the modeling of laser-plasma accelerators in a Lorentz boosted frame. The new implementation also enables the modeling of the generation of high-harmonic generation with intense lasers reflecting on a plasma mirror, with an accuracy that is out of reach with usual methods based on finite-difference electromagnetic solvers on existing supercomputers.

REFERENCES

- [1] I. Haber, R. Lee, H. Klein and J. Boris, *Advances In Electromagnetic Simulation Techniques*, in *Proc. Sixth Conf. Num. Sim. Plasmas*, Berkeley, CA, pp. 46-48, 1973.
- [2] J.-L. Vay, I. Haber and B. B. Godfrey, *Journal of Computational Physics* vol. 243, p. 260, Jun 2013.
- [3] B. Godfrey, *Journal of Computational Physics*, vol. 15, p. 504, 1974.
- [4] S. F. Martins, R. A. Fonseca, L. O. Silva, W. Lu and W. B. Mori, *Computer Physics Communications*, vol. 181, p. 869, May 2010.
- [5] J. L. Vay, C. G. R. Geddes, C. Benedetti, D. L. Bruhwiler, E. Cormier-Michel, B. M. Cowan, J. R. Cary and D. P. Grote, *AIP Conference Proceedings*, vol. 1299, p. 244, 2010.
- [6] J. L. Vay, C. G. R. Geddes, E. Cormier-Michel and D. P. Grote, *Journal of Computational Physics*, vol. 230, p. 5908, Jul. 2011.
- [7] L. Sironi and A. Spitkovsky, No Title, 2011.
- [8] B. B. Godfrey and J.-L. Vay, *Journal of Computational Physics*, vol. 248, p. 33, 2013.
- [9] X. Xu, P. Yu, S. F. Martins, F. S. Tsung, V. K. Decyk, J. Vieira, R. A. Fonseca, W. Lu, L. O. Silva and W. B. Mori, *Computer Physics Communications*, vol. 184, p. 2503, 2013.

- [10] B. Godfrey, *Journal of Computational Physics*, vol. 19, p. 58, 1975.
- [11] M. Kirchen, R. Lehe, B. B. Godfrey, I. Dornmair, S. J alas, K. Peters, J.-L. Vay and A. R. Maier, arXiv:1608.00215, 2016.
- [12] R. Lehe, M. Kirchen, B. B. Godfrey, A. R. Maier and J.-L. Vay, arXiv:1608.00227, 2016.
- [13] J.-L. Vay, *Physical Review Letters*, vol. 98, p. 130405, 2007.
- [14] B. Godfrey, *The IPROP Three-Dimensional Beam Propagation Code*, Defense Technical Information Center, 1985.
- [15] A. F. Lifschitz, X. Davoine, E. Lefebvre, J. Faure, C. Rechatin and V. Malka, *Journal of Computational Physics*, vol. 228, p. 1803, 2009.
- [16] A. Davidson, A. Tableman, W. An, F. Tsung, W. Lu, J. Vieira, R. Fonseca, L. Silva and W. Mori, *Journal of Computational Physics*, vol. 281, p. 1063, 2015.
- [17] R. Lehe, M. Kirchen, I. A. Andriyash, B. B. Godfrey and J.-L. Vay, *Computer Physics Communications*, vol. 203, p. 66, 2016.
- [18] R. Lehe, M. Kirchen, I. a. Andriyash, B. B. Godfrey and J.-L. Vay, *arXiv.org physics.pl*, 1507.04790v1, 2015, <http://arxiv.org/abs/1507.04790> arXiv:1507.04790.
- [19] J.-L. Vay, D. P. Grote, R. H. Cohen and A. Friedman, *Computational Science and Discovery*, vol. 5, p. 014019, 2012.
- [20] H. Vincenti, R. Lehe, R. Sasanka and J.-L. Vay, *Computer Programs in Physics*, to be published, 2016.
- [21] <http://www.fft.w.org>
- [22] <http://www.netlib.org/blas>
- [23] <http://numba.pydata.org>
- [24] G. Blaclard, H. Vincenti, R. Lehe and J.-L. Vay, arXiv:1608.05739, 2016.