

Efficient Modeling of Laser-Plasma Accelerators with INF&RNO

C. Benedetti, C. B. Schroeder, E. Esarey, C. G. R. Geddes and W. P. Leemans

Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA

Abstract. The numerical modeling code INF&RNO (INtegrated Fluid & paRticle simulation cOde, pronounced "inferno") is presented. INF&RNO is an efficient 2D cylindrical code to model the interaction of a short laser pulse with an underdense plasma. The code is based on an envelope model for the laser while either a PIC or a fluid description can be used for the plasma. The effect of the laser pulse on the plasma is modeled with the time-averaged ponderomotive force. These and other features allow for a speedup of 2-4 orders of magnitude compared to standard full PIC simulations while still retaining physical fidelity. The code has been benchmarked against analytical solutions and 3D PIC simulations and here a set of validation tests together with a discussion of the performances are presented.

Keywords: laser-plasma accelerators, simulation, benchmarking

PACS: 52.38.Kd, 47.11.-j, 05.10.-a

INTRODUCTION

Detailed and reliable numerical modeling in 3D of a laser-plasma accelerator (LPA) [1, 2], where a short and intense laser pulse interacts with an underdense plasma over distances ranging from a few millimeters/centimeters (yielding $\sim 0.1/1$ GeV electron energy [3–5]) up to a meter (expected ~ 10 GeV electrons [6]), is a formidably challenging task. A 3D "full" (i.e., where we take into account the fastest time scale represented by the oscillations of the laser field) particle-in-cell (PIC) simulation requires $10^4 - 10^5$ CPU hours in today's supercomputers for a millimeter-scale plasma and $O(10^6)$ CPU hours for a centimeter-scale plasma. Since in general the computational complexity for a fixed accuracy in the results grows more than linearly with the number of time steps, we easily deduce that a simulation of a meter-scale plasma requires tens of millions of CPU hours and so becomes unfeasible with standard simulation tools [7]. However, simulations are required, since the physics involved in the laser-plasma interaction is highly nonlinear and, consequently, analytical solutions are lacking. Numerical modeling plays a central role in helping our understanding of the physics. Two solutions have been proposed to overcome this limitation and allow for the simulation of multi-GeV LPA stages: *i.* run the full PIC simulation in an optimal boosted Lorentz frame [8] instead of in the laboratory frame; *ii.* use reduced models. The first option is certainly attractive and strongly pursued by several groups. The advantage of running a simulation in a boosted Lorentz frame relies on the fact that, if backward propagating waves (e.g., Raman backscattering) can be neglected, and this is usually true given the phenomenology of LPAs, then it has been shown [8] that the unbalance between the maximum and minimum physical scales involved in a simulation, which contribute to set the computational complexity of the problem, is not invariant under Lorentz transformation. It turns out that in general the laboratory frame is not the optimal choice to run the simulation while running it in a boosted frame can considerably reduce the scale unbalance, shortening (also by several orders of magnitude) the simulation length. Applications, estimated computational speedups, and limitations of this technique are discussed in [8–12]. Codes based on reduced models on the other hand allow for a significant speedup compared to full PIC simulations either because of dimensionality reduction (e.g., 2D cylindrical instead of full 3D cartesian) or because of approximations in the physical description of the system (e.g., quasi-static instead of fully dynamic plasma response, ponderomotive approximation instead of full Lorentz force, etc.). Even if they may lack important elements of the physics (e.g., a quasi-static code can not describe self-injection), their use has been proven to be successful in several relevant scenarios [13–17].

The INF&RNO computational framework, currently under development at LBNL, is a 2D cylindrical (r-z) code that adopts an envelope model for the laser pulse and makes use of the ponderomotive force approximation to describe the interaction of the laser pulse with the plasma. The plasma can be modeled using either a PIC or a fluid description and its response is fully dynamic even though a quasi-static module will be available soon. Both PIC and fluid modalities are integrated in the same computational framework allowing for staged simulations (e.g., PIC-mode for injection and

fluid-mode for acceleration). It is also possible to load and track self consistently externally injected bunches. The ultimate goal of the INF&RNO project is obtaining a fast reduced code suitable for modeling the relevant features of a LPA producing ~ 10 GeV in a meter-scale plasma (e.g., BELLA [6]) where, for a given problem, it is possible to switch between several physical descriptions/levels of approximations in order to clearly identify in each situation the relevant physics involved. In this paper we provide an overview of the INF&RNO framework (first section) together with a set of validation tests (second section) and a discussion of the performance and the future developments of the code (conclusion and outlook).

THE CODE: NUMERICS AND FEATURES

INF&RNO is a 2D-cylindrical (r-z) code which adopts non-dimensional, "comoving" variables defined as $\xi = k_p(z - ct)$ (longitudinal) and $\rho = k_p r$ (transverse), where $k_p = \omega_p/c$, ω_p is the plasma frequency corresponding to the chosen reference density n_0 , and c is the speed of light. The time is also rescaled with $1/\omega_p$, that is $\tau = \omega_p t$. The laser pulse is described using an envelope model [18]. Denoting by $a_\perp = eA_\perp/mc^2$ the normalized vector potential of the laser, the (slowly varying) envelope $\hat{a}$ is defined by $a_\perp = \frac{\hat{a}(\xi, \rho)}{2} e^{i(k_0/k_p)\xi} + c.c.$ The envelope evolves according to

$$\left(i \frac{k_0}{k_p} + \partial_\xi\right) \partial_\tau \hat{a} = \frac{1}{2} \frac{\delta}{\gamma_{\text{fluid}}} \hat{a} - \frac{1}{2} \nabla_\perp^2 \hat{a}, \quad (1)$$

where $2\pi/k_0$ is the central laser wavelength, $\delta = n/n_0$ is the (normalized) plasma density and γ_{fluid} is the relativistic factor associated with the local plasma fluid velocity (see below). The fully electromagnetic wakefield is described by the fields E_z, E_r, B_ϕ normalized to $E_0 = mc\omega_p/e$, where m and e are respectively mass and charge of the electron. The wakefield evolves according to Ampère-Maxwell laws which read

$$\frac{\partial E_r}{\partial \tau} = \frac{\partial(E_r - B_\phi)}{\partial \xi} - j_r, \quad \frac{\partial E_z}{\partial \tau} = \frac{\partial E_z}{\partial \xi} + \frac{1}{\rho} \frac{\partial(\rho B_\phi)}{\partial \rho} - j_z, \quad \frac{\partial B_\phi}{\partial \tau} = -\frac{\partial(E_r - B_\phi)}{\partial \xi} + \frac{\partial E_z}{\partial \rho}, \quad (2)$$

where (j_r, j_z) are the components of the (normalized) current density. The background plasma can be modeled using either a PIC or a fluid description while for external injected bunches only the PIC description is currently available. Laser-matter coupling is described via the ponderomotive approximation. The evolution equations for the PIC and fluid modalities are

$$\text{PIC} \rightarrow \begin{cases} \forall j=1, \dots, N_p \\ \frac{d\xi_j}{d\tau} = \frac{u_{z,j}}{\gamma_j} - 1 & \frac{du_{z,j}}{d\tau} = -\frac{\partial \gamma_j}{\partial \xi} - E_z - \frac{u_{r,j}}{\gamma_j} B_\phi \\ \frac{d\rho_j}{d\tau} = \frac{u_{r,j}}{\gamma_j} & \frac{du_{r,j}}{d\tau} = -\frac{\partial \gamma_j}{\partial \rho} - E_r + \frac{u_{z,j}}{\gamma_j} B_\phi \\ \gamma_j \equiv \sqrt{1 + |\hat{a}|^2/2 + u_{z,j}^2 + u_{r,j}^2} \end{cases}, \quad \text{fluid} \rightarrow \begin{cases} \frac{\partial \delta}{\partial \tau} = \frac{\partial \delta}{\partial \xi} - \nabla \cdot \left(\frac{\mathbf{u}}{\gamma_{\text{fluid}}} \delta \right) \\ \frac{\partial(\delta u_j)}{\partial \tau} = \frac{\partial(\delta u_j)}{\partial \xi} - \nabla \cdot (\tilde{\beta} \delta u_j) + \delta \left[-(\mathbf{E} + \frac{\mathbf{u}}{\gamma_{\text{fluid}}} \times \mathbf{B}) - \frac{1}{2\gamma_{\text{fluid}}} \nabla |\hat{a}|^2 \right]_j, \quad j=z, r \\ \gamma_{\text{fluid}} \equiv \sqrt{1 + |\hat{a}|^2/2 + u_z^2 + u_r^2}. \end{cases} \quad (3)$$

In the PIC part $(\xi_j, \rho_j, u_{z,j}, u_{r,j})$ are the phase-space coordinates (position and normalized momentum) of the j^{th} numerical particle representing one of the characteristics of the Vlasov equation for the plasma. In the fluid part the fields δ and $\mathbf{u} = (u_z, u_r)$ are the plasma density and momentum. The two modalities are integrated in the same computational framework, enabling an easy switch from one description to the other (combined simulations). A 1D "full" fluid version of the code is also available.

Concerning numerical aspects, all the fields are discretized into the same 2D mesh (no staggering is adopted). Longitudinal derivatives are computed using a second-order finite difference upwind scheme [19]: $(\partial_\xi f)_{i,j} = (-3f_{i,j} + 4f_{i+1,j} - f_{i+2,j})/(2\Delta_\xi)$, where $f_{i,j}$ is the field value at the (i, j) node and Δ_ξ the longitudinal cell size. Radial derivatives are computed using a standard centered second-order accurate scheme. The possibility to adopt a non-uniform radial grid (stretched radial grid) is currently under testing. No singularity exists at the $r = 0$ boundary, and from symmetry properties we have $\partial_\rho E_z|_{\rho=0} = E_r|_{\rho=0} = B_\phi|_{\rho=0} = 0$ and $\lim_{\rho \rightarrow 0} B_\phi/\rho = \partial B_\phi/\partial \rho|_{\rho=0}$. Second and fourth order Runge-Kutta integrators (RK2/RK4) are available for field evolution while plasma particles and externally injected bunches can be pushed with either RK4 or the standard Boris pusher [7]. The implementation of an implicit integration scheme is underway and is related to the inclusion of a quasi-static module for the wakefield and the plasma. Concerning force interpolation and charge/current deposition, both linear and quadratic shape functions have been implemented. Compact low-pass filters [20] are available for current and field smoothing. The user has large freedom




FIGURE 1. *Left:* particle splitting procedure for a "heavy" particle (the blue one) approaching the $r = 0$ axis. *Right:* longitudinal wakefield with and without splitting. Black: without splitting, 50 particles/cell and linear shape function. Red: without splitting, 50 particles/cell and quadratic shape function. Blue: with splitting (4 fragments), 9 particles/cell and quadratic shape function.




FIGURE 2. 1D fluid-envelope simulation: $a_0 = 1.5, k_0/k_p = 20, k_p L = 2$. *Left:* black plots are the laser envelope profile at $\omega_p t = 0$ and $\omega_p t = 1500$, the red/blue plots are the real/imaginary part of $\hat{a}$ at $\omega_p t = 1500$. *Right:* laser energy evolution deep into depletion as a function of the longitudinal resolution in the envelope (colored plots). The black plot is a full fluid simulation (no envelope approximation) at very high resolution (180 points/ λ_{laser}) and can be considered an "exact" result.

in loading numerical particles over the computational domain (the numerical particle distribution is controlled by a simple user-defined routine) and this freedom can be used to selectively provide a better sampling of the plasma phase space distribution within the dynamically interesting zones without greatly increasing the overall number of simulated particles. Because of the cylindrical symmetry, particles loaded at large radii carry generally more charge than particles loaded on-axis. If/when these "heavy" particles approach the $r = 0$ axis, they may induce "spikes" in density and currents increasing the noise level in the fields. In INF&RNO, to partially compensate this detrimental effect, particles with high charge approaching the axis are split into smaller fragments (see Fig. 1 (left)). The method is simple but quite effective as can be seen in Fig. 1 (right). Drawbacks of this approach are a small violation of the local charge/energy conservation and a local heating of the plasma while total charge and momentum are conserved.

We end this section with a comment on the validity of the envelope description when the laser is strongly depleted. Within the envelope approximation there is no need, in principle, to resolve the laser wavelength since the envelope length (L_{env}) is the smallest relevant scale as far as the pulse is concerned. However, during nonlinear laser evolution (redshifting, depletion) structures *smaller* than L_{env} arise (see Fig. 2 (left)) and the mesh resolution must be high enough to capture them. Fig. 2 (right) shows the evolution, deep into depletion, of the laser energy computed using the envelope model at various longitudinal resolutions (see figure caption for details). As long as the resolution is high enough, the envelope model correctly reproduces the physics still ensuring an effective speed-up (larger than ~ 2) compared to a "full" description of the laser. The computational savings from implementation of the envelope model will be reduced if modeling deep into depletion is required.

VALIDATION TESTS AND BENCHMARKS

Test 1. We consider the diffraction in vacuum (up to 10 Rayleigh lengths) of a tightly focused Gaussian pulse ($k_0 w_0 = 20, k_0 L = 40$, where w_0 and L are respectively the pulse waist and length): the goal is to test the accuracy of the laser envelope solver. In Fig. 3 (left) we show the evolution of the laser vector potential, $a(z)/a(z=0)$, as a




FIGURE 3. Left: vacuum diffraction of a focused laser pulse (see subsection "Test 1" for details). Center: time reversibility test for the vacuum diffracting pulse changing the resolution (see subsection "Test 1"). Right: mismatch oscillations of a low intensity pulse in a plasma channel (see subsection "Test 2" for details).




FIGURE 4. Comparison of the longitudinal wakefield obtained using INF&RNO with the 1D quasistatic nonlinear theory (left) and with VORPAL/OSIRIS results [24] (center, right). See subsection "Test 3" for details.

function of the propagation distance: the black curve is the analytic calculation, while the red and blue symbols are simulation results with two different resolutions (see figure). Time integration is performed with RK2. Simulations, even at low resolution, agree with theory. We also checked time reversibility of the simulation. After the forward propagation we evolved back in time the pulse measuring the error in recovering the initial condition. Denoting by $\hat{a}_{i,j}$ the (discretized) initial condition and by $\hat{a}'_{i,j}$ the recovered one, the reversibility error can be quantified according to $\epsilon_1 = \max_{i,j} \frac{|\hat{a}_{i,j} - \hat{a}'_{i,j}|}{a_0}$, or $\epsilon_2 = \sqrt{\frac{1}{N_z N_r} \sum_{i,j} \frac{|\hat{a}_{i,j} - \hat{a}'_{i,j}|^2}{a_0^2}}$. Results are shown in Fig. 3 (center) where we plot ϵ_1, ϵ_2 as a function of the resolution quantified by $w_0/\Delta r$ (we assumed $\Delta \xi = \Delta r$ and $\Delta t/\Delta \xi = 0.24$); a third order convergence is observed.

Test 2. The accuracy of the laser-plasma coupling has been checked by considering the evolution of a low intensity Gaussian pulse ($a_0 = 0.05, k_p L = 2, k_p w_0 = 2.5$) in a plasma channel ($k_0/k_p = 20, \Delta n = \Delta n_c$). The pulse is slightly mismatched ($\sim 5\%$) so we expect amplitude oscillations in the range $[0.05, 0.055125]$ with a period of $\omega_p T_{\text{osc}} = \pi k_p Z_{\text{Rayleigh}} = 196.3$. Simulation results showing the time evolution of a_0 are plot in Fig. 3 (right). The measured oscillation period is 196.6, in good agreement with the theoretical value. The damping of the oscillations, discussed in [21], is due to the fact that the (short) laser pulse is not monochromatic. Each chromatic component of the beam is characterized by a different oscillation frequency and the decoherence between these modes damps out the intensity oscillations. We can eliminate short pulse effects in the simulation removing the mixed derivative $\partial_{\xi, \tau}^2$ in the equation for the envelope evolution (1). Simulation results in this case (black dashed line in Fig. 3 (right)) show no damping and very good agreement with theory.

Test 3. We compare the longitudinal wakefield (E_z) obtained using INF&RNO with 1D analytical theory and with other codes (VORPAL[22], OSIRIS[23]) in the nonlinear regime. In Fig. 4 (left) we plot the on-axis lineout of E_z generated by a broad Gaussian laser pulse with $a_0 = 1, k_p L = 2, k_p w_0 = 8$. The black plot is the 1D quasistatic




FIGURE 5. Comparison between PIC and fluid modalities in INF&RNO: snapshots and lineouts for density and longitudinal wakefield (see subsection "Test 4" for details).

nonlinear theory while the red dashed line is the INF&RNO/fluid result ($k_p \Delta \xi = 1/20$, $k_p \Delta r = 1/10$, $\Delta t / \Delta z = 0.24$, RK2). The agreement in both phase and amplitude of the wake is very good. In Fig. 4 (center, right) we compare INF&RNO/PIC with the codes VORPAL and OSIRIS for $a_0 = 2, 4$. All the details concerning the benchmarking exercise (laser/plasma parameters, numerical settings for VORPAL and OSIRIS, etc.) can be found in [24]. Concerning INF&RNO/PIC the simulation has been done with $k_p \Delta \xi = 1/30$, $k_p \Delta r = 1/10$, $\Delta t / \Delta \xi = 0.24$, 20 particles/cell and quadratic shape function. The agreement is excellent for $a_0 = 2$ and good for $a_0 = 4$. We notice that in the $a_0 = 4$ case all the three codes slightly disagree in the rear part of the wake, but this is almost surely due to a numerical convergence issue.

Test 4. In this test we check the internal consistency between the PIC and the fluid modalities of INF&RNO. For the study we chose to operate in a mildly nonlinear regime where the cold fluid description for the plasma is still valid and, as a consequence, PIC and fluid calculations *must* give the same answer. The physical parameter for the laser pulse are $a_0 = 2.5$, $k_p L = 2$, $k_p w_0 = 5$. For the PIC simulation the numerical parameters are $k_p \Delta \xi = 1/30$, $k_p \Delta r = 1/20$, $\Delta t / \Delta \xi = 0.24$, 20 particles/cell (quadratic shape function). For the fluid one $k_p \Delta \xi = 1/25$, $k_p \Delta r = 1/10$, $\Delta t / \Delta \xi = 0.2$. Results are shown in Fig. 5 where we plot snapshots of the electron density and of the longitudinal wakefield from PIC and fluid runs. A comparison of the lineouts along $r = 0$ is also shown. Excellent agreement is observed.




FIGURE 6. Comparison between a full LPA simulation performed with the 3D PIC code ALADyn and INF&RNO (see subsection "Test 5" for details).

Test 5. In the last test presented we compare a full LPA simulation performed with the 3D-PIC code ALADyn [25, 26] with a run performed with INF&RNO. Details on the 3D ALADyn simulation are discussed in [27]: a (Gaussian) laser pulse ($a_0 = 5$, $\tau_{\text{fwhm}} = 30$ fs, $w_0 = 16$ μm) interacts with a 4 mm gas-jet (plasma density $3 \cdot 10^{18}$ e/cm³). The numerical parameters for the INF&RNO simulation are: $k_p \Delta \xi = 1/30$, $k_p \Delta r = 1/20$, $\Delta t / \Delta \xi = 0.25$, 6 particles/cell with particle

splitting to reduce on-axis noise and quadratic shape function. Results are presented in Fig. 6 where we compare the snapshots of the electron density at two different times for the two codes, the lineout of the accelerating field and the longitudinal phase space near the end of the simulation. The agreement is very good. The two codes capture basically the same physics (nonlinear evolution of the laser field, excitation of a bubble-like wake and injection) with the advantage that INF&RNO is more than 150 times faster compared to the 3D run.

CONCLUSION AND OUTLOOK

We have presented the INF&RNO computational framework, a 2D cylindrical, envelope, ponderomotive, PIC/fluid simulation code. The code has been widely tested and some of the benchmarks have been discussed in this paper. The plan for improving/upgrading the code foresees parallelization, enforcement of exact charge conservation, implementation of a quasi-static module and introduction of a boosted Lorentz frame modeling capability. Several performance profiling tests confirm a speed-up of several orders of magnitude (~ 2 -4 depending on the particular problem) compared to standard simulation tools. In particular, in the fluid modality, the CPU time required for one time step has been measured to be $\sim 0.6 \mu\text{s}/(\text{grid point})$. In this case, assuming reasonable resolution ($k_p \Delta \xi \sim 1/50$, $k_p \Delta r \sim 1/8$, $\omega_p \Delta t \sim k_p \Delta \xi / 4$) and computational domain size ($k_p L_z \sim k_p L_r \sim 20 - 30$), a LPA interaction over a meter-scale plasma (e.g., BELLA) will take less than 500 CPU hours, making feasible this kind of runs in a few days on small machines.

ACKNOWLEDGMENTS

This research is supported by the Department of Energy under the contract No. DE-AC02-05CH11231.

REFERENCES

1. T. Tajima, J.M. Dawson, Phys. Rev. Lett. **43**, p. 267 (1979).
2. E. Esarey, *et al.*, Rev. Mod. Phys. **81**, p. 1229 (2009).
3. S.P.D. Mangles, *et al.*, Nature **431**, p. 535 (2004); C.G.R. Geddes, *et al.*, *ibid.* p. 538; J. Faure, *et al.*, *ibid.* p. 541.
4. W.P. Leemans, *et al.*, Nature Physics **2**, p. 696 (2006).
5. K. Nakamura, *et al.*, Phys. Plasmas **14**, 056708 (2007).
6. W.P. Leemans, *et al.*, these Proceedings.
7. C.K. Birdsall, A.B. Langdon, Plasma Physics Via Computer Simulation, Adam Hilger, (1991).
8. J.-L. Vay, Phys. Rev. Lett. **98**, 130405 (2007).
9. S.F. Martins, *et al.*, in: Proc. 13th Advanced Accelerator Workshop, Santa Cruz, CA, p. 285 (2008).
10. D.L. Bruhwiler, *et al.*, in: Proc. 13th Advanced Accelerator Workshop, Santa Cruz, CA, p. 29 (2008).
11. S.F. Martins, *et al.*, Nature Physics **6**, p. 311 (2010).
12. J.-L. Vay, *et al.*, these Proceedings.
13. P. Mora, T.M. Antonsen, Phys. Plasmas **4**, p. 217 (1997).
14. C. Huang, *et al.*, J. Comp. Phys. **217**, p. 658 (2006).
15. N.H. Matlis, *et al.*, Nature Physics **2**, p. 749 (2006).
16. I. Blumenfeld, *et al.*, Nature **445**, p. 741 (2007).
17. A.F. Lifshitz, *et al.*, J. Comp. Phys. **228**, p. 1803 (2009).
18. E. Esarey, *et al.*, Phys. Fluids B **5**, p. 2690 (1993).
19. B.A. Shadwick, *et al.*, IEEE-TPS **30**, p. 38 (2002).
20. J.S. Shang, J. Comp. Phys. **153**, p. 312 (1999).
21. E. Esarey, W. P. Leemans, Phys. Rev. E **59**, p. 1082 (1999).
22. C. Nieter, J.R. Cary, J. Comp. Phys. **196**, p. 448 (2004).
23. R.A. Hemker, *et al.*, Lecture Notes in Computational Science **2331**, p. 342 (2002).
24. K. Paul, *et al.*, in: Proc. 13th Advanced Accelerator Workshop, Santa Cruz, CA, p. 315 (2008).
25. C. Benedetti, *et al.*, IEEE-TPS **36**, p. 1790 (2008).
26. C. Benedetti, *et al.*, Nuclear Inst. and Methods in Physics Research A **608**, p. S94 (2009).
27. C. Benedetti, in: Proc. 2nd Conf. on Ultraintense Laser Interaction Science, Frascati, Italy, AIP Conf. Proc. **1209**, p. 11 (2010).

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.