EFFICIENT MODELING OF LASER-PLASMA ACCELERATORS USING THE PONDEROMOTIVE-BASED CODE INF&RNO*

C. Benedetti†, C.B. Schroeder, E. Esarey, W.P. Leemans, LBNL, Berkeley, CA 94720, USA

Abstract

The ponderomotive force approximation enables efficient modeling of laser-plasma accelerators. It allows simulation in cylindrical geometry which captures relevant 3D physics at 2D computational cost. INF&RNO (INtegrated Fluid & paRticle simulatioN cOde) is an efficient 2D cylindrical code based on the envelope model for the laser, a PIC or fluid description for the plasma, and the ponderomotive force approximation to describe the effect of the laser pulse on the plasma. These and other features, such as an improved laser envelope solver, a dynamical resampling of the phase space distribution to reduce on-axis noise, and a Lorentz boosted frame modeling capability, allow for a speedup of several orders of magnitude compared to standard (explicit) full PIC simulations while still retaining physical fidelity. The code has been benchmarked against analytical solutions and 3D PIC simulations. In this paper we report on the latest developments of the code, focusing in particular on the improved laser envelope solver, and we discuss its performance.

INTRODUCTION

Numerical modeling of a laser-plasma accelerator (LPA) [1], where a short and intense laser pulse interacts with an underdense plasma over distances ranging from a few millimeters/centimeters (yielding ~ 0.1/1 GeV electron energy [2, 3]) up to several tens of centimeters (as in the BErkeley Lab Laser Accelerator experiment (BELLA) [4] where $\sim$ 10 GeV electrons are expected), is a computationally challenging task. A 3D "full" (*i.e.*, where the fast oscillations of the plasma electrons in the laser field are taken into account) PIC simulation requires $10^4 - 10^5$ CPU hours using today's supercomputers for a millimeter-scale plasma and $\sim 10^6$ CPU hours for a centimeter-scale plasma. The simulation of a ~ 10 GeV stage as required by BELLA would necessitate several tens of millions of CPU hours and are practically impossible to perform with standard simulation tools and today's computational resources. However, numerical modeling plays a central role in the understanding and optimization of LPAs. Simulations are required since the physics of the laser-plasma interaction is highly nonlinear and, consequently, analytical solutions are lacking.

Two approaches have been proposed to overcome these limitations and allow for the simulation of multi-GeV LPA stages: *i.* use reduced models [5, 6, 7]; *ii.* run the simulation in an optimal Lorentz boosted frame (LBF) [8] instead of in the laboratory frame. Codes based on reduced models 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 Lorenz 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 [5, 9, 10, 11]. The use of an LBF has been strongly pursued by several groups [12, 13, 14, 15]. The advantage of running a simulation in an LBF relies on the fact that, if backward propagating waves (*e.g.*, Raman backscattering) can be neglected, and this is generally true for 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.

The INF&RNO computational framework [6, 16], is a 2D cylindrical $(r - z)$ code that adopts an envelope model for the laser pulse and makes use of the (time-averaged) 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. Both PIC and fluid modalities are integrated in the same computational framework allowing for staged simulations (*e.g.*, PIC for injection and fluid for acceleration). The adoption of the cylindrical geometry allows the description of 3D physics (laser evolution, electromagnetic field structure) at 2D computational cost. The code features an improved laser envelope solver which enables an accurate description of the laser pulse evolution deep into depletion even at a reasonably low resolution. This new algorithm overcomes some of the limitations reported in other implementations of the laser envelope solver [7]. For the PIC part, a dynamical resampling of the phase space distribution is implemented in order to reduce the on-axis noise which affects, in some cases, 2D axisymmetric codes when the computational grid is loaded with a constant number of particles per cell. Finally, an LBF modeling capability has been introduced within the (noiseless) fluid framework. The employment of the LBF by the user is transparent since

* Work supported by the Office of Science, Office of High Energy Physics, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

† cbenedetti@lbl.gov

ISBN 978-3-95450-116-8

a set of "wrapper" routines take care of all the necessary data transformations between the laboratory frame and the LBF, where the simulation is performed, during initialization and output dumps. The code has been validated and benchmarked against analytical solutions and other codes (*e.g.*, fully 3D PIC [13]). Benchmark results can be found in [6]. In this paper we provide a general overview of the features and performance of the code. With INF&RNO, detailed simulations of a ∼10 GeV LPA stage becomes feasible in a few hours or days, depending on the particular simulation settings, on small clusters with (at most) a few hundred CPUs.

OVERVIEW OF THE CODE

Physical Model and Basic Equations

The code INF&RNO works in 2D cylindrical $(r-z)$ geometry and 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 $\tau = \omega_p t$. Working in 2D cylindrical geometry, unlike the 2D Cartesian, has the advantage that laser evolution (*e.g.*, self-focusing and diffraction) or wakefield structure are the same as in 3D, if non-axisymmetric features can be neglected, but the computational cost for the simulations is significantly lower (approximately 2 orders of magnitudes) compared to 3D runs.

Denoting by $a_\perp = eA_\perp/mc^2$ the normalized vector potential of the laser, the corresponding laser 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 [1]

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

where $2\pi/k_0$ is the central laser wavelength, $\delta = n/n_0$ is the (normalized) electron plasma density and γ_{fluid} is the relativistic factor associated with the local plasma fluid velocity. We notice that in the derivation of Eq. 1 all high-frequency contributions from the plasma have been neglected [1].

The fully time-explicit electromagnetic wakefield generated behind the laser 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

$$\begin{aligned}
\frac{\partial E_z}{\partial\tau} &= \frac{\partial E_z}{\partial\xi} + \frac{1}{\rho}\frac{\partial(\rho B_\phi)}{\partial\rho} - J_z,\\
\frac{\partial E_r}{\partial\tau} &= \frac{\partial(E_r - B_\phi)}{\partial\xi} - J_r, \qquad (2)\\
\frac{\partial B_\phi}{\partial\tau} &= -\frac{\partial(E_r - B_\phi)}{\partial\xi} + \frac{\partial E_z}{\partial\rho},
\end{aligned}$$

where (J_z, J_r) are the components of the (normalized) current density. The background plasma can be modeled using either a PIC or a fluid description and the laser-plasma coupling is described via the ponderomotive approximation [1]. For the PIC description the equations are

$$\begin{cases}
\frac{d\xi_j}{d\tau} = \frac{u_{z,j}}{\gamma_j} - 1 \equiv \beta_{z,j} - 1\\
\frac{d\rho_j}{d\tau} = \frac{u_{r,j}}{\gamma_j} \equiv \beta_{r,j}\\
\frac{du_{z,j}}{d\tau} = -\frac{\partial\gamma_j}{\partial\xi}|_j - E_{z,j} - \beta_{r,j}B_{\phi,j}\\
\frac{du_{r,j}}{d\tau} = -\frac{\partial\gamma_j}{\partial\rho}|_j - E_{r,j} + \beta_{z,j}B_{\phi,j}\\
\gamma_j \equiv (1 + |\hat{a}|^2/2 + u_{z,j}^2 + u_{r,j}^2)^{1/2}\,,
\end{cases} \qquad (3)$$

where $(\xi_j, \rho_j, u_{z,j}, u_{r,j})$ 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. A PIC description is also adopted to model externally injected particle beams. In the fluid description for the plasma, the electron density, δ, and momentum, $\mathbf{u} = (u_z, u_r)$, evolve according to

$$\begin{cases}
\frac{\partial\delta}{\partial\tau} = \frac{\partial\delta}{\partial\xi} - \nabla\cdot\left(\frac{\mathbf{u}}{\gamma_{\text{fluid}}}\delta\right)\\
\frac{\partial(\delta u_j)}{\partial\tau} = \frac{\partial(\delta u_j)}{\partial\xi} - \nabla\cdot\left(\vec{\beta}\delta u_j\right)\\
\quad + \delta\left[-(\mathbf{E} + \frac{\mathbf{u}}{\gamma_{\text{fluid}}}\times\mathbf{B}) - \frac{1}{2\gamma_{\text{fluid}}}\nabla\frac{|\hat{a}|^2}{2}\right]_j,\ j = z, r\\
\gamma_{\text{fluid}} \equiv \sqrt{1 + |\hat{a}|^2/2 + u_z^2 + u_r^2}\,.
\end{cases} \qquad (4)$$

The PIC and fluid modalities are integrated in the same computational framework, enabling an easy switch from one description to the other (combined simulations).

Recently, besides the time-explicit modality, a quasi-static, linear, fluid modality has been included in the INF&RNO framework. The quasi-static approximation assumes $\partial/\partial\tau = 0$ in the equations for the wakefield and the plasma, explicit time evolution is only retained in the equations for the driver (particle or laser beam). This modality will allow for fast simulations of the interaction of a (long) particle beam in an overdense plasma [*i.e.*, where the beam density is (much) smaller than the background plasma density] [17]. Finally, the code also features a 1D fluid modality (which can be either time-explicit or quasi-static) that can be used for quick/preliminary parameter scan.

Numerical Aspects and Features of the Code

In INF&RNO 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, namely $(\partial_\xi f)_{i,j} = (-3f_{i,j} + 4f_{i+1,j} - f_{i+2,j})/(2\Delta_\xi)$, 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. No singularity exists at the $r = 0$ boundary, and from symmetry properties we have, for instance, $\partial_\rho E_z|_{\rho=0} = E_r|_{\rho=0} = B_\phi|_{\rho=0} = 0$ and $\lim_{\rho\to 0} B_\phi/\rho = \partial B_\phi/\partial\rho|_{\rho=0}$. Second and fourth order Runge-Kutta integrators (RK2/RK4) are available for fluid plasma quantities and wakefield evolution while plasma particles and externally injected bunches can be pushed with either RK4 or the standard Boris pusher [18].

ISBN 978-3-95450-116-8

The laser envelope is updated in time solving Eq. 1 using a second order Crank-Nicolson scheme [16]. More details on the laser envelope solver are given in the next section.

In the PIC modality force interpolation, charge and current deposition are performed using quadratic shape functions. Compact low-pass digital filters [19] and standard binomial filters (with compensator) are available for field and current smoothing. The user has large freedom 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 [20]. When the computational grid is loaded with a constant number of particles per cell, particles loaded at large radii, because of the cylindrical symmetry, carry generally more charge than particles loaded on-axis. If/when these "heavier" particles approach the $r = 0$ axis, they may induce "spikes" in density and currents, increasing the noise level in the fields. These detrimental effects can be partially mitigated via dynamical particle splitting of the high charge particles approaching the axis. Details on the implementation of this technique to resample the phase space distribution are discussed in [6].

An LBF modeling capability is available within the INF&RNO/fluid framework. The extension of this capability to the INF&RNO/PIC modality is currently underway. In order to enable LBF simulations in a code based on a laser envelope model, the evolution equation for the laser envelope has to be Lorentz invariant. In the LHS of Eq. 1 we retain the second order time derivative (full wave operator), ensuring exact invariance of this equation under Lorentz transformation [16]. This simulation modality is transparent for the user which only has to set the velocity of the boosted frame (*i.e.*, γ_{LBF}) and initialize the system as he/she would be in the standard laboratory frame (LF). A set of "wrapper" functions automatically perform all the necessary operations to set properly the physical parameter [laser, plasma, externally injected bunch(es), *etc.*] and the numerical ones (grid settings) in the LBF. The swiping plane technique [15] is used to initialize the simulation in the LBF and to reconstruct the output data in the LF at a fixed laboratory time. In all the tests performed to validate the implementation of the LBF in the fluid framework, no evidence of the numerical instability which usually affects PIC simulations in the LBF when γ_{LBF} is large has been observed [21]. The theoretical speedup for an LPA simulation in the LBF and the optimal boost velocity ($\gamma_{\text{LBF}}^{opt}$) have been computed in [15]. It is found that, for a typical LPA experiment, $\gamma_{\text{LBF}}^{opt} \simeq \gamma_w$, where γ_w is the relativistic factor of the wake traveling behind the laser driver. A $\sim$10 GeV stage requires $n_0 \sim 10^{17}$ e/cm^3, which implies $\gamma_{\text{LBF}}^{opt} \simeq \gamma_w \sim 100$. However, in our case, the maximum value for γ_{LBF} (and so the maximum speedup) is more likely to be limited by the validity of Eq. 1 in the LBF. In fact, the envelope approximation relies on a time scale separation between the fast laser oscillations and the slowly evolving envelope. Due to the redshift of the laser light in the boosted frame, the fulfillment of the time scale separation condition becomes more difficult to achieve the higher is γ_{LBF}. We found that, for the $\sim$ 10 GeV stage, the maximum acceptable value for γ_{LBF} is around $30 - 40$. For higher values of γ_{LBF} the simulation results start to deviate significantly from those obtained in the LF.

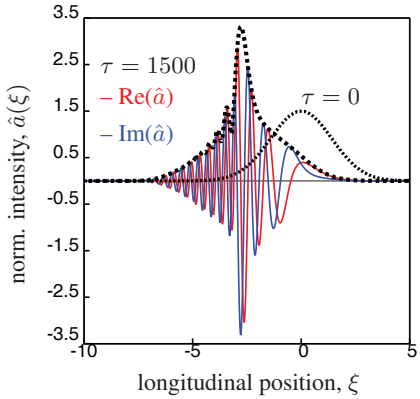

Figure 1: Evolution of the laser envelope in a 1D simulation with parameters $a_0 = 1.5, k_0/k_p = 20, L_{\text{rms}} = 1$ [$a^2(\xi, \tau = 0) = a_0^2 \exp(-\xi^2/2L_{\text{rms}}^2)$]. The dashed black lines refer to $\tau = 0$ (right) and $\tau = 1500$ (left). For $\tau = 1500$ then real (red) and imaginary (blue) part of $\hat{a}$ are also shown.

The code has been parallelized with MPI. Parallelization is achieved through 1D longitudinal domain decomposition. Studies to improve the parallel scalability are currently under consideration.

IMPROVED LASER ENVELOPE SOLVER

In a code based on a laser envelope model that adopts the ponderomotive approximation to describe the laser-matter interaction, there is no need, in principle, to resolve the laser wavelength and usually, for a resonant pulse [1], the characteristic length of the pulse (*e.g.*, the the r.m.s pulse length, L_{rms}) is the smallest relevant scale of interest. However, during propagation in the plasma, as a consequence of laser-pulse redshifting, structures smaller than L_{rms} arise in the laser envelope [22, 7, 23]. In Fig. 1, we show, as an example, the evolution of the laser envelope from $\tau = 0$ to $\tau = 1500$ in a 1D simulation with $a_0 = 1.5, k_0/k_p = 20, L_{\text{rms}} = 1$ [we assume, initially, a laser pulse of the form $a^2(\xi, \tau = 0) = a_0^2 \exp(-\xi^2/2L_{\text{rms}}^2)$]. The red and blue lines in Fig. 1 are, respectively, the real ($\Re[\hat{a}]$) and imaginary ($\Im[\hat{a}]$) part of the laser envelope at $\tau = 1500$. The presence, at later times, of structures smaller than L_{rms} (in the form of oscillations) in $\Re[\hat{a}]$ and $\Im[\hat{a}]$ (but also, to a minor extent, in the laser envelope $|\hat{a}|$) is evident. This has to be taken into account when designing a numerical scheme for the laser envelope evolution equation, or when choosing the (longitudinal) resolution to be used in the simulation.

ISBN 978-3-95450-116-8

In INF&RNO, the evolution equation for the laser envelope, Eq. 1, is discretized in time using a Crank-Nicholson scheme which reads

$$\begin{aligned} & - \quad \frac{\hat{a}^{n+1} - 2\hat{a}^{n} + \hat{a}^{n-1}}{\Delta_\tau^2} \\ & + \quad 2\left[i\frac{k_0}{k_p} + \frac{\partial}{\partial \xi} \right] \left(\frac{\hat{a}^{n+1} - \hat{a}^{n-1}}{2\Delta_\tau} \right) = \\ & - \quad \nabla_\perp^2 \left(\frac{\hat{a}^{n+1} + \hat{a}^{n-1}}{2} \right) + \frac{\delta^n}{\gamma^n_{\text{fluid}}} \frac{\hat{a}^{n+1} + \hat{a}^{n-1}}{2} , \end{aligned} \tag{5}$$

where $\hat{a}^k = \hat{a}(\xi, \rho; k\Delta_\tau)$. By using Eq. 5, knowing $\hat{a}^n$ and $\hat{a}^{n-1}$, we can evaluate $\hat{a}^{n+1}$.The first term in Eq. 5 corresponds to the (discretized) second order time derivative ($\partial^2/\partial\tau^2$) on the LHS of Eq. 1. Including this term, besides providing exact Lorentz invariance of the wave operator as required by the LBF modeling capability, it allows also for a correct description of backwards propagating waves and a better modeling of strongly depleted stages [23]. As for the spatial discretization, the transverse Laplacian operator, $\nabla_\perp^2$, takes the usual discrete form (for cylindrical coordinates) described in [18]. Extra care has to be put in the discrete representation of the longitudinal derivative $\partial/\partial\xi$. In fact, when the small structures in $\hat{a}$ develop and are not well resolved (this is the case for depleted laser pulses if the longitudinal resolution is not sufficiently high), a not optimal discrete form of the operator $\partial/\partial\xi$ might introduce significant numerical errors and prevent a correct description of the laser evolution. This issue is addressed in INF&RNO by means of a technique that involves the polar form for the representation of the complex field $\hat{a}(\xi)$, namely $\hat{a}^{(\text{polar})}(\xi) = a(\xi)\exp[i\theta(\xi)]$, where $a = |\hat{a}| = (\Re[\hat{a}]^2 + \Im[\hat{a}]^2)^{1/2}$ and $\theta = \arg(\hat{a})$, instead of the standard Cartesian representation. The polar amplitude and phase $a(\xi)$ and $\theta(\xi)$, are reasonably well behaved and less prone to show an oscillatory behavior or significant variability over small scales compared to $\Re[\hat{a}]$ or $\Im[\hat{a}]$. Evaluating the longitudinal derivative of the laser envelope field using the polar form has then some numerical advantage. More specifically, starting from the identity $\hat{a} = \hat{a}\, e^{-i\theta} e^{i\theta}$, and applying the $\partial/\partial\xi$ operator, we get

$$\frac{\partial \hat{a}}{\partial \xi} = \frac{\partial}{\partial \xi}\left(\hat{a}\, e^{-i\theta} e^{i\theta} \right) = \frac{\partial}{\partial \xi}\left(\hat{a} e^{-i\theta} \right) e^{i\theta} + i\hat{a}\frac{\partial \theta}{\partial \xi} , \tag{6}$$

where both $ae^{-i\theta} \equiv a(\xi)$ and $\theta(\xi)$ are, as previously described, reasonably smooth function, so we do not expect the numerical differentiation to introduce a significant error in the evaluation of their derivatives. We notice that in Eq. 5 the operator $\partial/\partial\xi$ acts on $\hat{a}^{n-1}$ and $\hat{a}^{n+1}$ (the unknown). To simplify the algorithm (retrieving the phase is a time consuming operation) we use the phase of $\hat{a}^n$ in evaluating both $\partial\hat{a}^{n+1}/\partial\xi$ and $\partial\hat{a}^{n-1}/\partial\xi$. The error associated with this approximation is small in all relevant cases analyzed.

An example (in 1D) of the performance of the laser envelope solver implemented in INF&RNO is shown in Fig. 2. In the simulation presented we consider the propagation of a short and intense laser pulse $a_0 = 1$, $L_{\text{rms}} = 1$ in a uniform plasma such that $k_0/k_p = 100$ (parameters of interest for a $\sim$ 10 GeV LPA stage). In Fig. 2 (a) we show the behavior of the laser energy as a function of the propagation distance (normalized to the pump depletion length which, for a resonant pulse, can be estimated as $L_{pd} \sim \lambda_p^3/\lambda_0^2 \sim 80$ cm, λ_p and λ_0 being respectively the plasma wavelength and the central laser wavelength) computed with different numerical schemes. The red dashed line is the result obtained with the full PIC code ALaDyn at high resolution, and can be considered "exact" for practical purposes. The black lines refer to results obtained with envelope code *without* the polar representation for the envelope field ("standard" solver). Different resolutions are considered, namely $L_{\text{rms}}/\Delta_\xi = 30, 100, 1000$ (see figure caption for detail). We notice that only at very high resolution the envelope results converge to the PIC result. The blue dots refer to the result computed with the envelope solver implemented in INF&RNO. Even at moderately low resolution, $L_{\text{rms}}/\Delta_\xi = 30$, the agreement with the PIC result is excellent also for a depleted laser pulse. Doubling the longitudinal resolution in this case (plot not shown) provides results in agreement with the PIC down to $L_{\text{propag}}/L_{pd} \simeq 1$. In Fig. 2 (b) we show lineouts of the longitudinal wakefield, E_z, after a propagation distance corresponding to 80% of the pump depletion length for the same numerical schemes presented before. Also in this case results obtained with the INF&RNO laser envelope model at (moderately) low resolution are in excellent agreement with the PIC simulation while, for the same resolution, results obtained with the "standard" solver show a significant (unacceptable) damping of the wakefield.

DISCUSSION

The computational framework INF&RNO is a 2D cylindrical, envelope, ponderomotive, PIC/fluid code designed to efficiently model LPAs. It features, among the others, an improved laser envelope solver, a technique to dynamically resample the phase space distribution aimed at reducing on-axis noise in the fields, and an LBF modeling capability. Compared to standard (explicit) full 3D PIC simulations it allows for a speedup of several orders of magnitude (between 2 and 5, depending on the particular problem and numerical setting) in the calculation time. In [16] it has been shown that an INF&RNO/fluid simulation of a 10 GeV stage in the quasi-linear regime, as in the BELLA experiment, requires only $\sim$ 30 CPU hours using the LBF modality with $\gamma_{\text{LBF}} = 12$. This corresponds to a speedup of approximately 5 orders of magnitude in the simulation time compared to standard simulation tools. As a second example we consider an INF&RNO/PIC simulation (in the laboratory frame) of a $\sim$ 10 GeV stage in the bubble regime where we model both self-injection and acceleration of the electron beam. Compared to the quasi-linear design the acceleration distance is shorter in this case ($\sim$ 10 cm for the bubble regime, $\sim$ 80 cm for the quasi-linear regime) but the

ISBN 978-3-95450-116-8

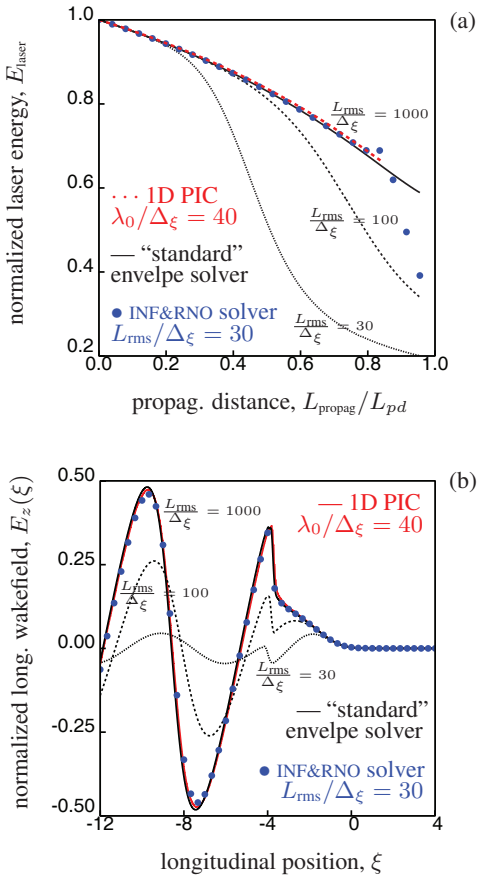

Figure 2: (a) Laser energy evolution as a function of the propagation distance normalized to the pump depletion length. (b) Lineout of the longitudinal wakefield, E_z, after a propagation distance corresponding to 80% of the pump depletion length. the simulation parameters are $a_0 = 1$, $L_{\rm rms} = 1$ and $k_0/k_p = 100$. In both (a) and (b) the red dashed line is the result obtained with the full PIC code ALaDyn. The black lines are the results obtained with the "standard" envelope solver at different resolution, namely $L_{\rm rms}/\Delta_\xi = 30, 100, 1000$. The blue points are the result obtained with the INF&RNO envelope solver for $L_{\rm rms}/\Delta_\xi = 30$.

background density is slightly higher ($\sim 3\times10^{17}$ e/cm^3 for the bubble regime, $\sim 1\times10^{17}$ e/cm^3 for the quasi-linear regime) and also the accelerating fields are higher (because of the higher background density and laser intensity) so the final beam energy is approximately the same. In order to correctly describe/represent the self-injection process and the nonlinear field structure of the bubble wake an adequate number of particles per cell and mesh resolution have to be adopted. In our test we chose $\Delta_\xi = 1/150$, $\Delta_r = 1/30$, and $\Delta_\tau/\Delta_\xi = 0.24$ with 12 particles per cell in the inner part of the computational domain (which includes the collection region for the trapped particles) and 2 particles per cell in the outer part. The size (normalized to k_p) of the computational box was ~ 20 (longitudinal) $\times$ 25 (radial). The simulation ran for ~ 5 days on 800 CPUs and the total cost was then $\sim 10^5$ CPU hours which corresponds to a speedup of about 2 orders of magnitude compared to standard tools. We notice that, even without an LBF modeling capability, accurate simulations of 10 GeV stages become feasible in a few days on relatively small computer clusters using INF&RNO.

ACKNOWLEDGEMENTS

The authors thank F. Rossi, L. Yu, J.-L. Vay, C.G.R. Geddes, and S. Rykovanov for insightful discussions.

REFERENCES

[1] E. Esarey, *et al.*, Rev. Mod. Phys. **81** (2009) 1229.

[2] S.P.D. Mangles, *et al.*, Nature **431** (2004) 535 (2004); C.G.R. Geddes, *et al.*, *ibid.* 538; J. Faure, *et al.*, *ibid.* 541.

[3] W.P. Leemans, *et al.*, Nature Physics **2** (2006) 696.

[4] W.P. Leemans, *et al.*, in: Proc. 2010 AAC Workshop, Annapolis, MD, p. 3, AIP (2010).

[5] P. Mora, T.M. Antonsen, Phys. Plasmas **4**, (1997) 217.

[6] C. Benedetti, *et al.*, in: Proc. 2010 AAC Workshop, Annapolis, MD, p. 250, AIP (2010).

[7] B.M. Cowan, *et al.*, J. Comp. Phys. **230** (2011) 61.

[8] J.-L. Vay, Phys. Rev. Lett. **98** (2007) 130405.

[9] A.J. Gonsalves, *et al.*, Nature Physics **7** (2011) 862.

[10] C. Huang, *et al.*, J. Comp. Phys. **217** (2006) 658.

[11] A.F. Lifshitz, *et al.*, J. Comp. Phys. **228** (2009) 1803.

[12] D.L. Bruhwiler, *et al.*, in: Proc. 2008 AAC Workshop, Santa Cruz, CA, p. 29 (2008).

[13] C. Benedetti, *et al.*, IEEE-TPS **36** (2008) 1790.

[14] S.F. Martins, *et al.*, Nature Physics **6** (2010) 311.

[15] J.-L. Vay, *et al.*, in: Proc. 2010 AAC Workshop, Annapolis, MD, p. 244, AIP (2010).

[16] C. Benedetti, *et al.*, Proc. of PAC2011, New York, MOP082 (2011), p. 250.

[17] C.B. Schroeder, *et al.*, Phys. Rev. Lett. **107** (2011) 145002.

[18] C.K. Birdsall, A.B. Langdon, Plasma Physics Via Computer Simulation, Adam Hilger, (1991).

[19] J.S. Shang, J. Comp. Phys. **153** (1999) 312.

[20] C. Benedetti, *et al.*, Nucl. Instr. and Meth. A **608** (2009) 94.

[21] J.-L. Vay, *et al.*, J. Comp. Phys. **230** (2011) 59083.

[22] C.B. Schroeder, *et al.*, Phys. Rev. Lett. **106** (2011) 135002.

[23] W. Zhu, *et al.*, Phys. Plasmas **19**, (2012) 033105.

ISBN 978-3-95450-116-8