Fast multidimensional model for the simulation of Raman amplification in plasma

J. P. Farmer and A. Pukhov
Heinrich Heine Universität, 40215 Düsseldorf, Germany
(Received 17 October 2013; published 10 December 2013)

We present Leap, a simulation model for Raman amplification in plasma, combining an envelope treatment of the laser fields with an electrostatic particle-in-cell solver. The code is fully two dimensional, with the model readily extendible to three dimensions, and includes dispersive and refractive effects. Simulations carried out for Raman amplification in a plasma channel show that guiding of both the pump and the probe contribute to the evolution of the probe, resulting in a shorter, more intense pulse.

DOI: 10.1103/PhysRevE.88.063104 PACS number(s): 52.38.Bv, 52.40.Fd, 52.65.−y

I. INTRODUCTION

Raman amplification in plasma has been suggested as a possible method to realize the next generation of ultrashort, ultraintense laser pulses [1]. As plasma supports fluences several orders of magnitude greater than conventional amplification media, it greatly reduces the size of the system by reducing or even avoiding the need to compress the amplified pulse. "Damage" to the gain medium is not a concern, as the excited plasma wave used to couple the pump and probe laser pulses is itself generated by their interaction.

As such, Raman amplification has been the focus of many experimental [2–5] and theoretical [6–10] works. However, for simulation of parameters relevant to experiments, there is typically some tradeoff between the computational overhead and the completeness of the model. Particle-in-cell (PIC) models offer the most complete treatment [10], but require significant computational resources. Three-wave models are the fastest [6], but omit particle effects such as wave breaking. Envelope-PIC models have been successfully used in other areas of laser-plasma interactions [11,12] and have been developed for the treatment of Raman amplification [7,13], with the goal of including most of the relevant physics with a significantly lower computational overhead than full PIC codes. Here we present the development of the Leap model—an envelope-PIC treatment which builds on previous works by including multidimensional effects, dispersion, and refraction in the context of Raman amplification.

The paper is arranged as follows: Section II discusses the laser solver, with Sec. III discussing the electrostatic solver. The two are coupled through the plasma response, detailed in Sec. IV. A case study investigating amplification in a plasma channel is made in Sec. V to illustrate the applicability of the code, with conclusions drawn in Sec. VI. The numerical treatment of the laser and electrostatic solvers are discussed in Appendixes A and B, respectively.

II. LASER SOLVER

We start with the wave equation, transforming into a moving frame: $\tau = t$, $\xi = z/c - t$. We limit our analysis to a single transverse dimension y using a slab geometry. (Extension of the model to three dimensions is, however, straightforward.) By normalizing the transverse coordinate by a factor $1/c$, we obtain

$$\left(\partial_\tau^2 - 2\partial_\tau\partial_\xi - \nabla_\perp^2\right)\vec{a} = \chi\vec{a}, \tag{1}$$

where $\vec{a}$ is the reduced vector potential of the laser field, $\vec{a} = e\vec{A}/mc$. We have made use of the quasistatic approximation for electron response, considering only the fast electron motion due to the laser field, with χ the normalized local plasma susceptibility, $\chi = \rho e/\varepsilon_0\gamma m$ (for a homogeneous, nonrelativistic plasma, $\chi = -\omega_p^2$). We substitute for separate pump and probe envelopes $\vec{a} = a_0\,\mathrm{e}^{i\phi_0}\,\vec{u} + a_1\,\mathrm{e}^{i\phi_1}\,\vec{u}$, where a_0, a_1 are, respectively, the pump and probe envelope amplitudes, which may be chosen complex to allow detuning from the envelope phases, $\phi_0 = w_0(\xi + 2\tau)$, $\phi_1 = -\omega_1\xi$, which satisfy the vacuum dispersion relation. $\vec{u}$ is the polarization vector, $\vec{u} = (\hat{x} + i\,\hat{y})/\sqrt{2}$ for circularly polarized light. We have neglected longitudinal components of the laser field, consistent for initial conditions satisfying $\vec{a}\cdot\hat{\xi} = 0$.

We separate for the pump and probe by Fourier analyzing and neglecting nonresonant terms:

$$\begin{aligned}\left(\partial_\tau^2 - 2\partial_\tau\partial_\xi + 2i\omega_0(\partial_\tau - 2\partial_\xi) - \nabla_\perp^2\right)a_0 &= \chi\left(a_0 + a_1\,\mathrm{e}^{i(\phi_1-\phi_0)}\right),\\ \left(\partial_\tau^2 - 2\partial_\tau\partial_\xi + 2i\omega_1\partial_\tau - \nabla_\perp^2\right)a_1 &= \chi\left(a_1 + a_0\,\mathrm{e}^{i(\phi_0-\phi_1)}\right).\end{aligned} \tag{2}$$

In an envelope treatment, counterpropagating modes manifest as short-wavelength modulations of the envelope. In order to reduce the required resolution, and consistent with the slowly varying envelope approximation, the second derivative in τ is often suppressed, preventing backscatter and reflection while retaining the dispersive behavior of the laser pulse [11,12]. Here, the use of separate envelope solvers for the pump and probe allow this term to be neglected, lowering the required resolution, while retaining backscatter and reflection. We expect the influence of dispersion on the long pump pulse to be negligible, and so additionally neglect derivatives $\partial_\tau\partial_\xi a_0$ and self-currents χa_0 in the pump evolution equation, which simplifies the boundary conditions for the laser solver. Retaining these terms for the probe allows effects such as dispersion and Raman forward scattering to be modeled.

The resulting equations can be solved directly, but this requires a high resolution as χ varies on the scale of half a laser wavelength. This requirement for high resolution may be avoided by introducing new averaged variables, the average susceptibility $\tilde{\chi} = \langle\chi\rangle$, and the coupling susceptibility $\tilde{\psi} = \langle\chi\,\mathrm{e}^{i(\phi_0-\phi_1)}\rangle$:

$$\begin{aligned}(2i\omega_0(\partial_\tau - 2\partial_\xi) - \nabla_\perp^2)a_0 &= \tilde{\psi}^* a_1,\\ (-2\partial_\tau\partial_\xi + 2i\omega_1\partial_\tau - \nabla_\perp^2)a_1 &= \tilde{\chi}a_1 + \tilde{\psi}a_0.\end{aligned} \tag{3}$$

 ©2013 American Physical Society

The average susceptibility governs the refraction and dispersion of a single laser pulse. The coupling susceptibility is related to the mutual index of refraction as used by Shvets *et al*. [1], governing the scattering between the two pulses and electromagnetically induced guiding [14]. It provides a measure of the electron bunching and is equivalent to the plasma wave amplitude used in three-wave models for Raman scattering [6]. The calculation of $\tilde{\psi}$ is functionally similar to the approach used by Hur *et al*. [7] to calculate the currents, although the implementation is somewhat different.

The complete numerical implementation of the laser update is detailed in Appendix A.

III. ELECTROSTATIC SOLVER

Envelope-PIC models have been effectively used in the study of wakefield acceleration [11,12], removing the need to resolve the laser wavelength. Using a Yee solver with an additional ponderomotive term added to the particle push allows a much lower resolution to be used, greatly decreasing the computational overhead. However, for the case of Raman backscattering, the smallest spatial scale is that of the excited plasma wave. An envelope treatment of the laser fields is not sufficient to reduce the computational overhead, as a high spatial resolution is still required to model the plasma wave, which, for an explicit method such as the Yee algorithm, will in turn require a high temporal resolution.

Hur *et al*. [7,13] overcame this limitation by using an electrostatic solver for the plasma response. The laser envelopes were solved on a low-resolution grid, with each cell subdivided into smaller cells used to model the electrostatic field. This separation of scales allows a high resolution to be used for the electrostatic solver with a long time step.

We take a similar approach, using an electrostatic solver for the plasma response, the resolution of which may be varied independently of the laser solver. Treating the system as electrostatic is a reasonable approximation, as the excited plasma wave is electrostatic, and the number of trapped particles is necessarily small in regimes suitable for efficient amplification.

We start with Gauss' law for electrostatics (note the use of normalized coordinates):

$$\nabla^2\varphi(y,\xi) = -c^2\rho(y,\xi). \tag{4}$$

As we use a moving coordinate frame, the plasma wave will extend beyond the downstream window boundary. We therefore treat the electric field as periodic only in the transverse y plane. A sensible choice is then a Fourier and cyclic redundancy (FACR) method [15]. We take the Fourier transform in y to give

$$\left(\partial_\xi^2 - k_y^2\right)\hat{\varphi} = -c^2\hat{\rho}. \tag{5}$$

Appendix B discusses the numerical solution of this equation.

IV. PLASMA RESPONSE

Coupling between the laser and electrostatic solvers is achieved through the plasma particles, which respond to both ponderomotive and electrostatic forces. The electron positions then determine the charge density used in the electrostatic solver and the average and scattering susceptibilities used in the laser solver.

For the case of Raman scattering, we assume stationary ions. The electron momentum may be considered as the sum of a slowly varying component and the fast "quiver" momentum of the electron in the laser field. For the case of an electrostatic plasma wave, the slowly varying momentum is equal to the canonical momentum $\vec{P}$ [11] to give a total momentum $\vec{P} - \vec{a}mc$. The canonical momentum represents the electron response to the electrostatic and ponderomotive forces:

$$\begin{aligned}\frac{d\vec{P}}{d\tau} &= -e\vec{E} - \frac{m}{2\gamma}\nabla|\vec{a}|^2 \\ &= -e\vec{E} - \frac{m}{4\gamma}\big(\nabla|a_0|^2 + \nabla|a_1|^2 \\ &\quad - 2(\omega_0+\omega_1)\mathrm{Im}\big(a_0 a_1^* \,\mathrm{e}^{i(\phi_0-\phi_1)}\big)\hat{\xi}\big).\end{aligned} \tag{6}$$

The macroparticle position depends only on the canonical momentum, representing the equilibrium position of an electron oscillating in the laser field. The average Lorentz factor depends on the total momentum [12]:

$$\begin{aligned}\gamma^2 &= 1 + \frac{\vec{P}^2}{m^2c^2} + |\vec{a}|^2 \\ &= 1 + \frac{\vec{P}^2}{m^2c^2} + \frac{1}{2}\big(|a_0|^2 + |a_1|^2 + 2\mathrm{Re}\big(a_0 a_1^* \,\mathrm{e}^{i(\phi_0-\phi_1)}\big)\big).\end{aligned} \tag{7}$$

The charge density ρ and the average and coupling susceptibilities $\tilde{\chi}$ and $\tilde{\psi}$ are calculated by summing over all particles:

$$\begin{aligned}\rho(y,\xi) = \frac{1}{\delta y\delta\xi}\sum_j \int_{y-\delta y/2}^{y+\delta y/2}\int_{\xi-\delta\xi/2}^{\xi+\delta\xi/2} \\ f(y_j - y', \xi_j - \xi')q_j dy' d\xi',\end{aligned} \tag{8}$$

$$\begin{aligned}\tilde{\chi}(y,\xi) = \frac{1}{\delta y\delta\xi}\sum_j \int_{y-\delta y/2}^{y+\delta y/2}\int_{\xi-\delta\xi/2}^{\xi+\delta\xi/2} \\ f(y_j - y', \xi_j - \xi')\frac{q_j^2}{\varepsilon_0\gamma_j m_j} dy' d\xi',\end{aligned} \tag{9}$$

$$\begin{aligned}\tilde{\psi}(y,\xi) = \frac{1}{\delta y\delta\xi}\sum_j \int_{y-\delta y/2}^{y+\delta y/2}\int_{\xi-\delta\xi/2}^{\xi+\delta\xi/2} \\ f(y_j - y', \xi_j - \xi')\frac{q_j^2}{\varepsilon_0\gamma_j m_j} \,\mathrm{e}^{i\Delta\phi_j}\, dy' d\xi',\end{aligned} \tag{10}$$

where y_j denotes the y coordinate of the jth particle, etc., with $\Delta\phi_j = (\omega_0 - \omega_1)\,\xi + 2\omega_0\tau$ the difference between the pump and probe carrier phases, ϕ_0 and ϕ_1, at the position of the jth particle. $\delta y, \delta\xi$ are the cell sizes in y and ξ, and f is the particle shape function, each of which may be chosen independently for the electrostatic and laser solvers.

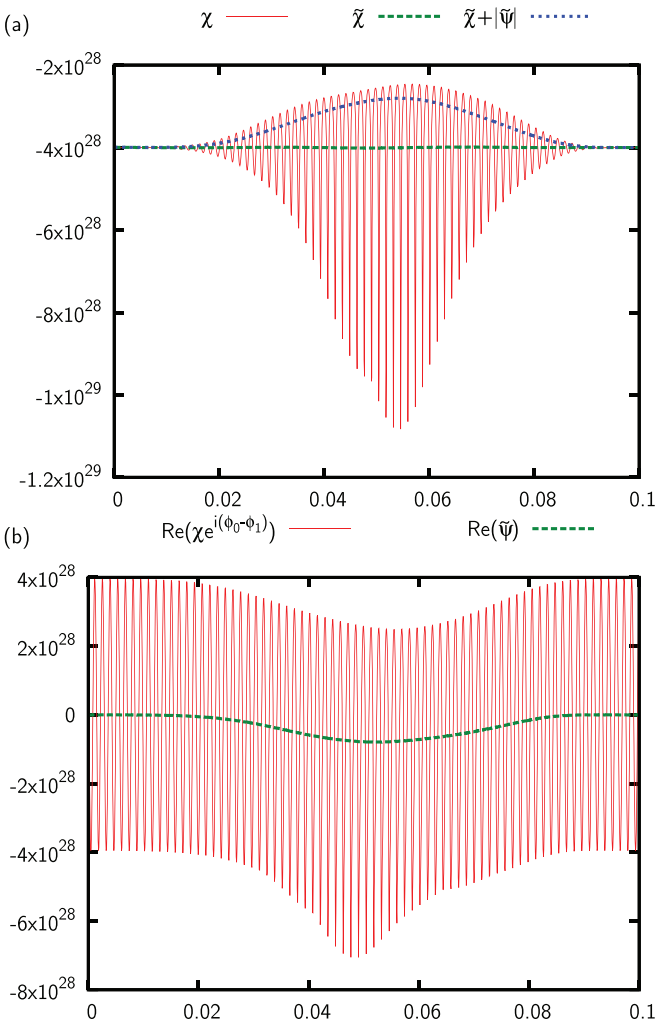

FIG. 1. (Color online) Plots showing the plasma susceptibility χ and the slowly varying quantities $\tilde{\chi}$ and $\tilde{\psi}$. (a) It can be seen that the slowly varying quantities $\tilde{\chi}$ and $\tilde{\psi}$ together represent the fully resolved susceptibility, χ. $\tilde{\chi}$ gives the average susceptibility, while the coupling susceptibility $\tilde{\psi}$ gives a measure of the bunching. (b) The coupling susceptibility is calculated as the average of the rapidly varying quantity $\chi\, e^{i(\phi_0-\phi_1)}$.

The particle shape function is an important aspect of the model, as it allows the rapidly varying quantities χ, $\chi\, e^{i(\phi_0-\phi_1)}$ to be averaged to their slowly varying equivalents $\tilde{\chi}$, $\tilde{\psi}$. We chose a second-order particle shape, with a length greater than the bunching length to ensure $\tilde{\chi}$, $\tilde{\psi}$ are smooth, as shown in Fig. 1.

Hur *et al.* [7] overcame the same problem by setting the cell size to equal the beat length $2\pi c/(\omega_0+\omega_1)$. Summing over a cell, the oscillatory behavior of $\chi\, e^{i(\phi_0-\phi_1)}$ is removed, leaving only the slowly varying component. This approach works well in the linear regime, in which the laser wavelengths are well known. However, our approach is more general, allowing the use of arbitrary resolution. This is necessary to allow the simulation of ultrashort pulses and has an advantage for chirped pulses, which lead to a varying beat length, and nonlinear regimes, in which evolution of the probe wavenumber means the beat length is not known.

V. RAMAN AMPLIFICATION IN A PLASMA CHANNEL

Previous reduced models for the simulation of Raman amplification range from simple three-wave models [6], which do not include particle effects such as wave breaking, to quasi-2D reduced PIC models [13] designed to investigate diffraction and gain focusing. The model presented here represents a significant advance, as it is fully two dimensional, and includes dispersion and refraction, allowing the simulation of effects such as self-modulation, guiding, self-focusing, and filamentation. Such effects are clearly of great importance if Raman amplification is to be used as an alternative to conventional amplification techniques: if a probe diffracts too rapidly, it ceases to be ultraintense, while if it filaments, efficient amplification will not be achieved and the resulting pulse will not be suitable for many practical applications.

We present a case study illustrating the capabilities of the code. Simulation parameters have been chosen to be similar to those used in experimental works [2,3]: a flat-top 800-nm pump pulse with a diameter of 55 μm and normalized amplitude $a_0 = 0.01$; a probe pulse with a FWHM (amplitude) duration of 500 ps and diameter of 55 μm with a normalized amplitude $a_1 = 0.001$; interacting in 2 mm of plasma of uniform density 1.3×10^{19} cm^{-3}. The probe frequency was chosen to be downshifted from the pump by the plasma frequency, corresponding to a wavelength of 876 nm.

The evolution of the probe pulse is shown in Fig. 2(a), with the pump amplitude, coupling susceptibility, and averaged susceptibility shown for the end of the interaction. The probe pulse is simultaneously amplified and compressed, from an amplitude $a_1 = 0.001$, duration 500 fs, to $a_1 = 0.056$, duration 270 fs. The plot of the coupling susceptibility shows the excited plasma wave breaks, which prevents further interaction between the pump and probe. The on-axis pump amplitude decreases by 50%, and so both pump depletion and wave breaking likely contribute to the compression of the probe. While estimates of the energy are, by necessity, limited for a planar geometry, an energy amplification of 1220 times was observed, corresponding to an energy transfer from pump to probe of 35%.

A simulation was then made for the case of amplification in a plasma channel. Such schemes have been demonstrated experimentally through the use of a plasma capillary [4], and through the creation of a plasma channel inside a gas jet [5]. We choose a parabolic density profile with the density at the simulation boundary 1.5 times that on axis, corresponding to an increase of 5.6% at the probe waist. As we expect the channel to guide the laser pulses, we set the pump diameter to the matched value for this channel, 18.5 μm. Note that the pump is introduced at the upstream simulation boundary, and so the evolution of the pump up to that point is not modeled. Full PIC simulations covering the entire interaction region have been carried out by Trines *et al.* [10]. The pump power is maintained, giving a pump amplitude of $a_0 = 0.017$. The results of the simulation are shown in Fig. 2(c). An increase may be seen in both pulse amplification and compression, with the amplified probe amplitude $a_1 = 0.106$ and duration 165 fs. Less pump depletion is observed than for the case of flat plasma, with an on-axis reduction of 40%, although the energy gain is somewhat higher, with an increase of 1360 times, corresponding to an energy-transfer efficiency of 40%.

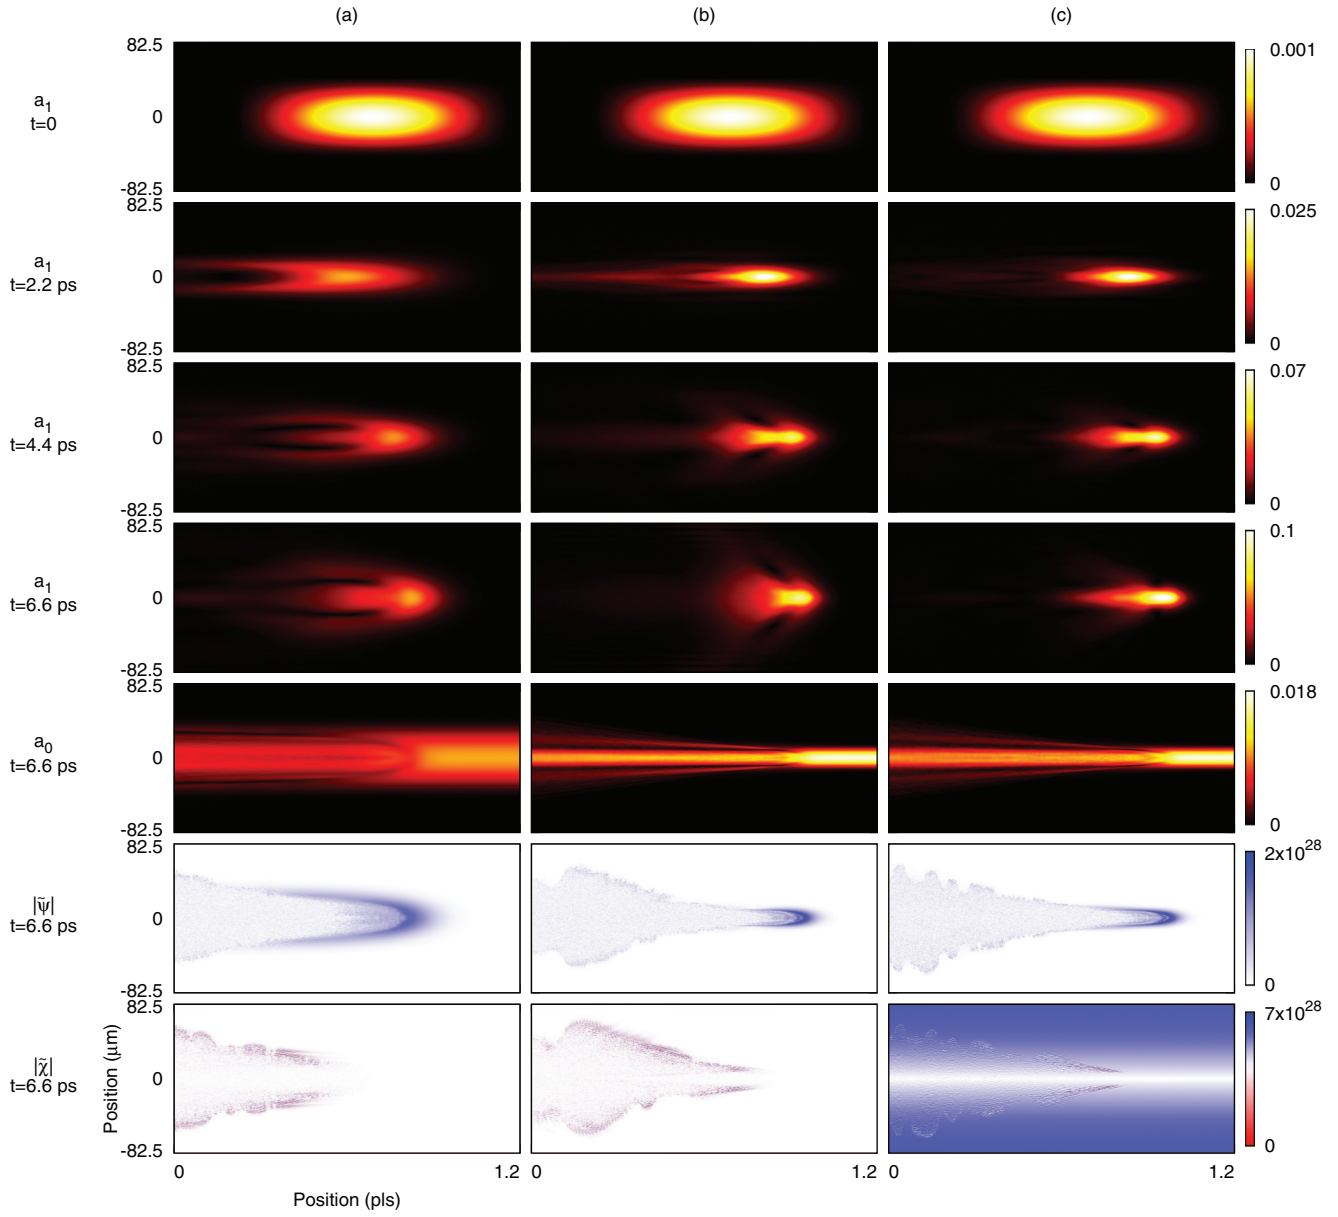

FIG. 2. (Color online) Plots showing the influence of a plasma channel on Raman amplification: (a) flat plasma, 55-μm pump, (b) flat plasma, 18.5-μm pump (nonphysical), and (c) plasma channel, 18.5-μm pump. The evolution of the probe is shown, along with the pump, coupling susceptibility and averaged susceptibility corresponding to the final probe snapshot. For comparison to typical experimental values, coordinates are given in micrometers (transverse) and picolightseconds (longitudinal).

In order to ascertain to what extent this marked difference in the probe evolution is due to refractive guiding of the probe, we consider a third case: an 18.5-μm pump in flat plasma. While this situation is nonphysical, as such a pump would diffract over the interaction length, its use in simulations allows us to weigh the relative influence of the narrow pump in simulations of the plasma channel. The results are shown in Fig. 2(b). It can be seen that this combination of parameters gives an intermediate result—amplification is greater than for the case of a 55-μm pump in flat plasma, but remains lower than the case of the 18.5-μm pump in a plasma channel. We may therefore conclude that refractive guiding of both the pump and probe contribute to the increase in amplitude observed in the case of a plasma channel.

VI. CONCLUSIONS

The development of Leap, a simulation model for the efficient simulation of Raman amplification in plasma, is discussed. It combines an envelope model for the laser pulses with an electrostatic PIC solver for particle motion, leading to a significant reduction in computational overhead. The discussion in this work is limited to a slab geometry, but the model is readily extensible to three dimensions. The envelope model includes dispersion and refraction, similar to those used in the study of wakefield acceleration [11,12], but with the significant addition that backscattering and reflection may also be modeled, achieved through the use of two counterpropagating laser envelopes and the addition of a

coupling term. The use of an electrostatic PIC solver to allow a separation of scales was first suggested by Hur [7], but this work further develops this idea through the generalization of the electrostatic solver to two dimensions, and the introduction of higher-order macroparticles to remove the limitations on the laser frequencies and simulation resolution.

The applicability of the code is demonstrated through the simulation of Raman amplification in a plasma channel. It is observed that the use of a plasma channel gives a more focused pulse with a higher peak amplitude and a shorter duration than amplification in a plasma of uniform density; an amplitude amplification of over 100 times was observed for the case of a plasma channel, and compression from 500 to 165 fs, compared to an amplification of 56 times and compression to 270 fs for flat plasma. It is shown that both refractive guiding of the probe and the narrower, more intense pump used in the plasma channel case contribute to this improvement.

The model opens the possibility to simulate parameter spaces relevant to experimental works, including the physics of multidimensional effects, without the prohibitive computational overhead associated with conventional simulation models.

ACKNOWLEDGMENT

This work was funded by DFG Transregio SFB TR18 and EU FP7 EuCARD^2.

APPENDIX A: NUMERICAL IMPLEMENTATION OF THE LASER UPDATE

We consider the evolution of the probe as given by Eq. (3). The full wave equation in the laboratory frame may be factorized to reveal two counterpropagating solutions, both of which exist for a single pulse propagating in a dispersive medium. Here, as we have neglected the second derivative in τ, the result is two solutions, one near-stationary in the probe frame, the other one near-stationary in the laboratory frame:

$$\partial_\tau(-2\partial_\xi + 2i\omega_1)a_1 = \tilde{\chi}a_1 + \nabla_\perp^2 a_1 + \tilde{\psi}a_0. \tag{A1}$$

The nonpropagating solution replaces the counterpropagating solution, suppressing backscatter and reflection while maintaining the dispersive behavior of the pulse. We introduce an auxiliary field α_1, allowing us to write

$$\partial_\tau a_1 = \alpha_1, \tag{A2}$$

$$(-2\partial_\xi + 2i\omega_1)\alpha_1 = \tilde{\chi}a_1 + \nabla_\perp^2 a_1 + \tilde{\psi}a_0. \tag{A3}$$

This scheme has the advantage that there is no numerical dispersion in vacuum for waves propagating along the ξ axis (as $\alpha_1 = 0$). We discretize these equations, using the notation $a|_m^n$ to denote quantity a at $\xi = m\,\delta\xi$, $\tau = n\,\delta\tau$, where $\delta\xi$ is the longitudinal grid spacing and $\delta\tau$ is the time step. For clarity, we have dropped the notation for the position in y, as the only transverse dependence is due to the transverse Laplacian, $\nabla_\perp^2$.

We solve these mutually dependent equations using a leapfrog-in-time method, evaluating the field at time step n, $n+1$, and the auxiliary field at $n+1/2$, $n+3/2$, such that

$$\partial_\tau a_1\big|_m^{n+1/2} = \alpha_1\big|_m^{n+1/2}, \tag{A4}$$

$$(-2\partial_\xi + 2i\omega_1)\alpha_1\big|_{m+1/2}^{n+1/2} = (\tilde{\chi}a_1 + \nabla_\perp^2 a_1 + \tilde{\psi}a_0)\big|_{m+1/2}^{n+1/2}. \tag{A5}$$

Substituting derivatives for their Taylor expansions and interpolating for values between grid points, we obtain

$$\frac{1}{\delta\tau}\left(a_1\big|_m^{n+1} - a\big|_m^n\right) = \alpha_1\big|_m^{n+1/2} \tag{A6}$$

$$\begin{aligned}
&\frac{2}{\delta\xi}\left(\alpha_1\big|_m^{n+1/2} - \alpha_1\big|_{m+1}^{n+1/2}\right) + i\omega_1\left(\alpha_1\big|_m^{n+1/2} + \alpha_1\big|_{m+1}^{n+1/2}\right)\\
&\quad= \frac{1}{4}\left(\tilde{\chi}\big|_m^{n+1}a_1\big|_m^n + \tilde{\chi}\big|_m^{n+1}a_1\big|_m^{n+1}\right.\\
&\qquad\left. + \tilde{\chi}\big|_{m+1}^{n+1}a_1\big|_{m+1}^n + \tilde{\chi}\big|_{m+1}^{n+1}a_1\big|_{m+1}^{n+1}\right)\\
&\qquad+ \frac{1}{4}\left(\nabla_\perp^2 a_1\big|_m^n + \nabla_\perp^2 a_1\big|_m^{n+1} + \nabla_\perp^2 a_1\big|_{m+1}^n + \nabla_\perp^2 a_1\big|_{m+1}^{n+1}\right)\\
&\qquad+ \frac{1}{2}\left(\tilde{\psi}\big|_m^{n+1}a_0\big|_m^n + \tilde{\psi}\big|_{m+1}^{n+1}a_0\big|_{m+1}^n\right).
\end{aligned} \tag{A7}$$

Note that we have assumed $\tilde{\chi}|_m^n \approx \tilde{\chi}|_m^{n+1}$, i.e., the average susceptibility evolves slowly in the probe frame. This value reflects the combined laser wake of the pump and probe; the intensity of both the short probe and the long pump evolve slowly in this frame, and so we can expect the same of the associated wake terms. This reduces the memory overhead, as the previous values of $\tilde{\chi}$ do not need to be stored. We also limit the interpolation of $\tilde{\psi}a_0$ to two points, as small errors in the growth rate have very little influence on the probe evolution, and so a lower degree of accuracy may be tolerated than in calculations of the probe dispersion, which may lead to significant nonphysical growth or damping if not correctly time centered. This reduces both the memory overhead, as the previous values of $\tilde{\psi}$ do not need to be stored, and the computational overhead, by allowing considerably simpler calculations.

We know the upstream value of α_1 is zero, as $a_1 = 0$, allowing us to integrate backwards from the right window boundary. We use Eq. (A6) to substitute for values of $a_1|_m^{n+1}$ in Eq. (A7), and rearrange to give

$$\begin{aligned}
&\left(1 + \frac{i\omega_1\delta\xi}{2} - \frac{\delta\tau\delta\xi}{8}\tilde{\chi}\big|_m^{n+1} - \frac{\delta\tau\delta\xi}{8}\nabla_\perp^2\right)\alpha_1\big|_m^{n+1/2}\\
&\quad= \left(1 - \frac{i\omega_1\delta\xi}{2} + \frac{\delta\tau\delta\xi}{8}\tilde{\chi}\big|_{m+1}^{n+1} + \frac{\delta\tau\delta\xi}{8}\nabla_\perp^2\right)\alpha_1\big|_{m+1}^{n+1/2}\\
&\qquad+ \frac{\delta\xi}{4}\left(\tilde{\chi}\big|_m^{n+1}a_1\big|_m^n + \tilde{\chi}\big|_{m+1}^{n+1}a_1\big|_{m+1}^n\right)\\
&\qquad+ \frac{\delta\xi}{4}\left(\nabla_\perp^2 a_1\big|_m^n + \nabla_\perp^2 a_1\big|_{m+1}^n\right)\\
&\qquad+ \frac{\delta\xi}{4}\left(\tilde{\psi}\big|_m^{n+1}a_0\big|_m^n + \tilde{\psi}\big|_{m+1}^{n+1}a_0\big|_{m+1}^n\right).
\end{aligned} \tag{A8}$$

The presence of the transverse Laplacian on the left-hand side requires that an implicit method be used, with an entire slice (constant ξ) evaluated simultaneously. While Eq. (3) is equivalent to that used by Cowan *et al.* [12], the introduction of an auxiliary field allows the use of simpler numerical techniques: in the two-dimensional case treated here, the

problem takes the form of a tridiagonal matrix and is solved using the Thompson algorithm.

APPENDIX B: ELECTROSTATIC SOLVER

Taking Eq. (5) and substituting for the derivative using Taylor expansions, we obtain

$$\hat{\varphi}|_{m+1} + \hat{\varphi}|_{m-1} = (2 + \delta\xi^2 k_y^2)\hat{\varphi}|_m - c^2\delta\xi^2\hat{\rho}|_m. \quad \text{(B1)}$$

If the upstream boundary conditions are known, this equation may be solved directly by integrating backward from the boundary. However, this is not the case here: while φ at $\xi = \infty$ may be taken as zero, this does not give the value of φ at the window boundary. We instead follow the method developed by Buneman [15]. The treatment here corresponds to the case of a purely imaginary ω in that work, allowing the equations to be recast in a simplified form.

Choosing $e^{\kappa} + e^{-\kappa} = 2 + \delta\xi^2 k_y^2$, i.e., $\cosh(\kappa) = 1 + \delta\xi^2 k_y^2/2$, we may factorize Eq. (B1) to give

$$(\hat{\varphi}|_{m+1} - e^{-\kappa}\,\hat{\varphi}|_m) - e^{\kappa}(\hat{\varphi}|_m - e^{-\kappa}\,\hat{\varphi}|_{m-1}) = -c^2\delta\xi^2\hat{\rho}|_m. \quad \text{(B2)}$$

Defining $\vartheta|_m = \hat{\varphi}|_m - e^{-\kappa}\,\hat{\varphi}|_{m-1}$, we can write

$$\vartheta|_{m+1} - e^{\kappa}\,\vartheta|_m = -c^2\delta\xi^2\hat{\rho}|_m. \quad \text{(B3)}$$

This is solved using a two-pass method, with ϑ calculated from right-to-left and $\hat{\varphi}$ calculated from left-to-right:

$$\vartheta|_m = e^{-\kappa}(\vartheta|_{m+1} + c^2\delta\xi^2\hat{\rho}|_m), \quad \text{(B4)}$$

$$\hat{\varphi}|_m = \vartheta|_m + e^{-\kappa}\,\hat{\varphi}|_{m-1}. \quad \text{(B5)}$$

As the charge density for all points beyond the upstream boundary is zero, it follows that ϑ at the upstream boundary is zero, allowing the right-to-left pass. However, the moving-window geometry results that plasma beyond the downstream boundary has been perturbed, and so the charge density is unknown. Despite this, from the quasineutral nature of plasma, we estimate φ at the downstream boundary to be zero. Any error in this boundary condition will have a diminishing influence as we advance the solution in ξ, as $|e^{-\kappa}| < 1$. This contrasts to the case of solving Eq. (B1) directly, where errors in the boundary condition can have an increasing influence while advancing in ξ (due to incorrect values of $\partial_\xi\varphi$). φ is then found by taking the inverse Fourier transform.

[1] G. Shvets, N. J. Fisch, A. Pukhov, and J. Meyer-ter-Vehn, Phys. Rev. Lett. **81**, 4879 (1998).

[2] W. Cheng, Y. Avitzour, Y. Ping, S. Suckewer, N. J. Fisch, M. S. Hur, and J. S. Wurtele, Phys. Rev. Lett. **94**, 045003 (2005).

[3] J. Ren, W. Cheng, S. Li, and S. Suckewer, Nat. Phys. **3**, 732 (2007).

[4] G. Vieux, A. Lyachev, X. Yang, B. Ersfeld, J. P. Farmer, E. Brunetti, R. C. Issac, G. Raj, G. H. Welsh, S. M. Wiggins, and D. A. Jaroszynski, New J. Phys. **13**, 063042 (2011).

[5] C.-H. Pai, M.-W. Lin, L.-C. Ha, S.-T. Huang, Y.-C. Tsou, H.-H. Chu, J.-Y. Lin, J. Wang, and S.-Y. Chen, Phys. Rev. Lett. **101**, 065005 (2008).

[6] V. M. Malkin, G. Shvets, and N. J. Fisch, Phys. Rev. Lett. **82**, 4448 (1999).

[7] M. S. Hur, G. Penn, J. S. Wurtele, and R. Lindberg, Phys. Plasmas **11**, 5204 (2004).

[8] N. A. Yampolsky, N. J. Fisch, V. M. Malkin, E. J. Valeo, R. Lindberg, J. Wurtele, J. Ren, S. Li, A. Morozov, and S. Suckewer, Phys. Plasmas **15**, 113104 (2008).

[9] T.-L. Wang, D. Michta, R. R. Lindberg, A. E. Charman, S. F. Martins, and J. S. Wurtele, Phys. Plasmas **16**, 123110 (2009).

[10] R. M. G. M. Trines, F. Fiuza, R. Bingham, R. A. Fonseca, L. O. Silva, R. A. Cairns, and P. A. Norreys, Nat. Phys. **7**, 87 (2011).

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

[12] B. M. Cowan, D. L. Bruhwiler, E. Cormier-Michel, E. Esarey, C. G. R. Geddes, P. Messmer, and K. M. Paul, J. Comput. Phys. **230**, 61 (2011).

[13] M. S. Hur and J. S. Wurtele, Comput. Phys. Commun. **180**, 651 (2009).

[14] G. Shvets and A. Pukhov, Phys. Rev. E **59**, 1033 (1999).

[15] O. Buneman, J. Comput. Phys. **12**, 124 (1973).