Tree-code simulations of proton acceleration from laser-irradiated wire targets

P. Gibbon[a)]
John von Neumann Institute for Computing, ZAM, Forschungszentrum Jülich GmbH, D-52425 Jülich, Germany

F. N. Beg
Department of Mechanical and Aerospace Engineering, University of California, San Diego, California 92093

E. L. Clark and R. G. Evans
AWE plc, Aldermaston, Reading RG 4PR, United Kingdom

M. Zepf
Department of Physics, The Queens University, Belfast BT7 1NN, United Kingdom

(Received 8 January 2004; accepted 4 May 2004; published online 19 July 2004)

Recent experiments using Terawatt lasers to accelerate protons deposited on thin wire targets are modeled with a new type of gridless plasma simulation code. In contrast to conventional mesh-based methods, this technique offers a unique capability in emulating the complex geometry and open-ended boundary conditions characteristic of contemporary experimental conditions. Comparisons of ion acceleration are made between the tree code and standard particle-in-cell simulations for a typical collisionless ''hole boring'' scenario in slab geometry. The utility of the gridless approach is emphasized by a series of simulations in ''wire'' geometry, in which electrons are permitted to circulate around the target at arbitrary distances from the focal region. The simulations reveal a number of features in common with recent experimental observations, including a disclike emission pattern of the MeV protons accelerated away from the wire. © *2004 American Institute of Physics.* [DOI: 10.1063/1.1767096]

I. INTRODUCTION

Since the first experiments measuring ion emission from multi-Terawatt laser–solid interactions,[1,2] laser-induced acceleration of MeV protons (fast ions) has become one of the most contentious issues in the field. Such protons originate from water vapor or other impurities adsorbed onto the target surface prior to laser irradiation, and by virtue of their lower mass, are preferentially accelerated over heavier constituent plasma ions when the laser creates a charge separation either inside or outside the target. The ability to create multi-MeV protons in a relatively cheap and compact manner has generated widespread interest because of its potential in a number of emerging fields, such as hadron therapy,[3] novel neutron sources,[4] and advanced fusion concepts.[5] Experimental campaigns begun by the Livermore and Imperial College groups some four years ago resulted in two apparently irreconcilable pictures of proton acceleration.[6–8]

The first interpretation, originally put forward by the Livermore team,[6,9] supposes that protons will be primarily accelerated from the *rear* surface of thin (1–100 μm) foil targets by the space charge setup by the laser-generated hot electron cloud. This intuitive scenario, dubbed ''target normal sheath acceleration,'' or TNSA, has since been strongly supported by two- and three-dimensional particle-in-cell (PIC) simulations performed by various authors over the last three years.[9–11] These simulations, based on a self-consistent solution of the Lorentz–Maxwell equations for the electromagnetic fields and plasma electrons and ions, all show an efficient initial transfer of laser energy to MeV electrons, which proceed virtually unhindered through the target and beyond. A large charge separation is thus rapidly created on the rear side, which then tugs ions away from this surface.

An alternative school of thought argues that most of energetic protons in high intensity interactions must come from the *front* side of the target, a viewpoint supported by experiments performed by the Imperial College London group at the Rutherford–Appleton Laboratory[12–14] and by the Michigan group.[15] The details of the mechanism for the ''front-side'' scenario are still unclear, however, ponderomotively driven charge separation and the associated ion shock formation appears—according to PIC simulations—to be insufficient by itself to account for the high number, energies, and angular distribution of protons observed. Only at intensities in excess of $10^{20}\,\mathrm{W\,cm^{-2}}$ and target thicknesses below around 10 μm is the front-side mechanism predicted to deliver higher ion energies than the rear-side blowoff.[16]

It is generally acknowledged that both mechanisms probably play a role, the real bone of contention is which one dominates for a particular laser–target configuration. In order to probe the physics of proton acceleration further, recent campaigns by the London and Darmstadt groups have been carried out using different target geometries.[17,18] In particular, a series of experiments with the VULCAN laser using *wire* targets has added fuel to this debate, as well as throwing up new questions concerning the role of ''spectator'' targets

[a)]Electronic mail: p.gibbon@fz-juelich.de

which appear to radically alter the field distribution in the vicinity of the laser-irradiated region.[19]

The purpose of this paper is to report on simulations of ion acceleration from wire targets using the parallel tree code PEPC (Pretty Efficient Parallel Coulomb-solver). Like the particle-in-cell method, this technique also follows the motion of charged particles in self-consistent electric (and in principle magnetic) fields. In contrast to PIC, however, the tree code computes interparticle potentials and forces *directly* rather than by employing a grid to mediate the fields via charge and current densities. As will become apparent shortly, this mesh free, Lagrangian approach lends itself rather well to the kind of open-ended, complex geometry typical of contemporary high intensity laser–matter interactions. After an introductory description of the tree-code-based model in Secs. II–IV, simulations of proton acceleration from laser-irradiated wire targets are presented for parameters close to conditions in the recent Imperial College of London (ICL) Rutherford–Appleton Laboratory (RAL) experiments.

II. FINITE-SIZE PARTICLE KINETICS WITH A PARALLEL TREE CODE

The hierarchical tree method upon which PEPC is based actually has more in common with molecular dynamics than with particle-in-cell simulation. Briefly, this technique makes systematic use of multipole expansions to reduce the computational effort expended in the force summation to a time $O(N \log N)$, which for large systems of charges ($N>10^4$) leads to substantial speed ups over the conventional $O(N^2)$ algorithm, independently of machine architecture. The technical details of the parallel algorithm used have been documented elsewhere,[20] and we will concentrate on the main components of the physical model in what follows. An earlier plasma tree code (in many respects a sequential forerunner to PEPC) has previously been used to perform microscopic molecular dynamics (MD) simulations of dense, strongly coupled plasmas.[21]

In the laser–plasma context of interest here, we use the tree algorithm to model ''macroscopic'' plasma behavior in the same spirit as PIC or fluid simulation. This model is based on the ''finite-size-particle'' (FSP) approach, in which point particles are replaced by spherical clouds, and are allowed to interpenetrate or cross each other. A detailed theoretical basis for this approach was actually laid down over 30 years ago by Langdon, Okuda, and Birdsall.[22,23] An important outcome of their work was to show that the collisionality of FSP plasmas is reduced by orders of magnitude compared to a plasma comprising point particles, so that the plasma parameter $n\lambda_D^3$ is in some sense replaced by $n\varepsilon^3$, where n, λ_D are the number density and Debye length, respectively, and ε is a measure of the particle size, or cloud radius. This property is implicitly and deliberately exploited in PIC codes, where the smoothing arises automatically by the imposition of a spatial grid, with the result that the above parameters are typically restricted to values, $\varepsilon \simeq \Delta x \simeq \lambda_D$. Henceforth, we will use the term FSP to mean *gridless* particle simulation.

The pure FSP method has two immediate advantages over PIC in kinetic plasma simulation.

(i) Collisions are included in a more controllable manner through the choice of ε/a, where $a=n^{-1/3}$ is the average interparticle spacing, and do not need to be patched back into the code in an *ad hoc* (and usually expensive) fashion.[24]

(ii) There are no geometrical restrictions on the simulation region; fast (laser-accelerated) particles do not have to be artifically absorbed or recycled, and may fly as far as they wish away from the interaction region. This does not preclude the imposition of periodic or reflective boundary conditions for special geometries; a fully periodic system for strongly coupled plasmas was developed, for example, in Ref. 25.

The drawback of the model is that it is, for the time-being, purely electrostatic; induced magnetic fields are neglected and no electromagnetic wave propagation is supported. At first sight, this may seem too simplistic to describe the kind of highly relativistic, nonlinear phenomena which prevail in high-energy-density laser–matter interactions. As we shall see, however, this ansatz does in fact allow us to capture the salient features of ion acceleration, including important collisional physics which has evidently been missing from the vast majority of PIC simulations of laser–solid interactions to date.

We now proceed with a ''formal'' description of the electrostatic FSP model as currently implemented in PEPC. The choice of units is somewhat subtle for macroscopic mesh-free plasma simulation, and contrasts with the microscopic ''Debye'' system used, for example in Ref. 21. The base normalizations for time, space, velocity, charge, and mass, respectively, are as follows:

$$\begin{aligned} t &= \omega_p^{-1}\tilde{t}, \\ v &= c\tilde{v}, \\ \mathbf{r} &= c\,\omega_p^{-1}\tilde{\mathbf{r}}, \\ q &= N_p e\tilde{q}, \\ m &= N_p m_e \tilde{m}. \end{aligned} \tag{1}$$

The constant N_p represents the number of physical charges contained within a simulation (macro-) particle, to be determined through the equation of motion, which for a given particle i with charge q_i and mass m_i is given (in cgs units) by

$$m_i \frac{d\mathbf{u}_i}{dt} = q_i \mathbf{E}_j = q_i \sum_{i \neq j} \frac{q_j \mathbf{r}_{ij}}{r_{ij}^3}, \tag{2}$$

where $\mathbf{r}_{ij}=\mathbf{r}_i-\mathbf{r}_j$ is the separation between particles i and j, and $\mathbf{u}_i=\gamma\mathbf{v}_i$ is its proper velocity; $\gamma=(1+|\mathbf{u}|^2/c^2)^{1/2}$ the relativistic factor. In a tree code, the $O(N)$ sum over all other particles is replaced by a sum over *multipole* expansions (expanded here up to quadrupole) of groups of particles, whose size increases with distance from particle i. The number of terms in this sum is $O(\ln N)$, which even after the additional overhead in computing the multipoles results in a substantial saving in effort for large N (Ref. 26).

Rewriting Eq. (2) in terms of the normalized variables (1), we find

$$\tilde{m}_i \frac{d\tilde{\mathbf{u}}_i}{dt} = \frac{N_p e^2 \omega_p}{m_e c^3} \tilde{q}_i \sum_{i \neq j} \frac{\tilde{q}_j \tilde{\mathbf{r}}_{ij}}{\tilde{r}_{ij}^3},$$

which after adding an external field $\mathbf{E}^p$ and making use of the plasma frequency definition, $\omega_p^2 = 4\pi e^2 n_e / m_e$ for electron density n_e, reduces to

$$\tilde{m}_i \frac{d\tilde{\mathbf{u}}_i}{dt} = \frac{1}{3} \tilde{q}_i \sum_{i \neq j} \frac{\tilde{q}_j \tilde{\mathbf{r}}_{ij}}{\tilde{r}_{ij}^3} + \tilde{q}_i \mathbf{E}^p(\mathbf{r}_i), \tag{3}$$

provided we take

$$N_p = \frac{4\pi}{3} n_e \left(\frac{c}{\omega_p} \right)^3. \tag{4}$$

Physically, the constant N_p is just the number of electrons in a sphere with radius c/ω_p. Since it has been normalized out, we do not actually need to know N_p in order to carry out a simulation, although it does provide a convenient conversion factor.

As in classical MD simulation, we cannot use the pure Coulomb law for point charges because of the finite time step, which will cause some particles to experience large, stochastic jumps in their acceleration, eventually destroying the energy conservation. We therefore modify the force law in Eq. (2) to include a softening parameter ε, so that the electric field looks like

$$\mathbf{E}(\mathbf{r}) = \frac{q\mathbf{r}}{(r^2 + \varepsilon^2)^{3/2}}. \tag{5}$$

The effect of the softening parameter is to introduce a cutoff into the potential, and to ensure that $\mathbf{E}(\mathbf{r}) \to 0$ as $r \to 0$, which greatly assists numerical stability in the time-integration (or particle-pusher) scheme. Physically, we no longer have point charges, but rather charge clouds with a smooth charge density. It is instructive to compute the latter by applying Gauss' law to Eq. (5), giving

$$\rho(r) = \frac{3q\varepsilon^2}{4\pi(r^2 + \varepsilon^2)^{5/2}}. \tag{6}$$

Using the same normalizations as before, and taking $\rho = e n_0 \tilde{\rho}$, where n_0 is some number density to be determined, we find

$$n_0 \tilde{\rho}(\tilde{r}) = n_e \frac{\tilde{q}\tilde{\varepsilon}^2}{(\tilde{r}^2 + \tilde{\varepsilon}^2)^{5/2}}. \tag{7}$$

To simplify this expression, we choose $n_0 = n_e$, or $\tilde{n}_e = 1$. Charge assignment is then straightforward: the total charge contained within a cuboid volume $V = x_L \times y_L \times z_L$ (in normalized units) is

$$Q = \sum_i q_i = \tilde{\rho}_0 V = N_e Q_s,$$

where N_e is the total number of simulation electrons and Q_s is the macrocharge carried by them. Since the initial density $\tilde{\rho}_0 = -\tilde{n}_e = -1$, we simply have

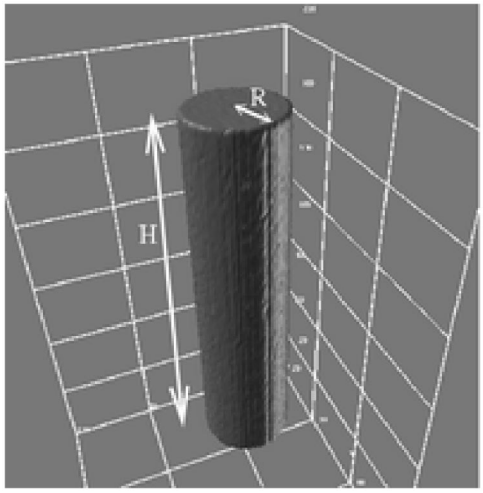

FIG. 1. Geometry for laser–wire simulations.

$$Q_s = -\frac{V}{N_e}. \tag{8}$$

III. TARGET PREPARATION

Assigning charges Q_s and $-Q_s Z$ to the electrons and ions, respectively, and masses $M_s^e = |Q_s|$, $M_s^i = A|Q_s|$, where Z and A are the atomic number and mass, sets up a macroscopic plasma system whose internal dynamics is governed solely by Eq. (3). Before we can proceed, however, we must pay some attention to its initial spatial and thermal configuration. Whereas a PIC code can be fairly easily initialized through a ''quiet start''—an orderly placement of particles in phase space—the FSP model suffers the same kind of pitfalls encountered in classical MD simulation, such as: (i) strong initial heating resulting from the system being out of equilibrium at $t = 0$, and/or (ii) persistent drift currents and oscillations due to localized random concentrations of ion charge.

In the present work, these problems are resolved by a two-step ''target preparation'' phase. First, ions are forced into a quasicrystalline structure bounded by the target geometry (which could be, for example: cuboid, wedge-shaped, or cylindrical). This is efficiently achieved by allowing the ions to interact via an artificial Lennard-Jones-type potential (the Coulomb interaction having been switched off), thus collectively seeking out a spatial configuration such that the mean distance to each nearest neighbor is maximized. The wire targets of the present investigation are constructed from cylinders of length H and radius R, as depicted in Fig. 1. The laser is focused either at the midpoint along the z axis or with some offset z_0.

Next, electrons are placed close to the ions (assuming $Z = 1$) with a velocity randomly selected from a Maxwellian distribution with temperature T_e. The whole system is then allowed to relax with the Coulomb force-law reinstated and with the additional thermodynamic constraint that $T_e = \text{const}$. (Ref. 21). This allows the system to seek out its own minimum potential energy while maintaining the temperature desired for the actual simulation.

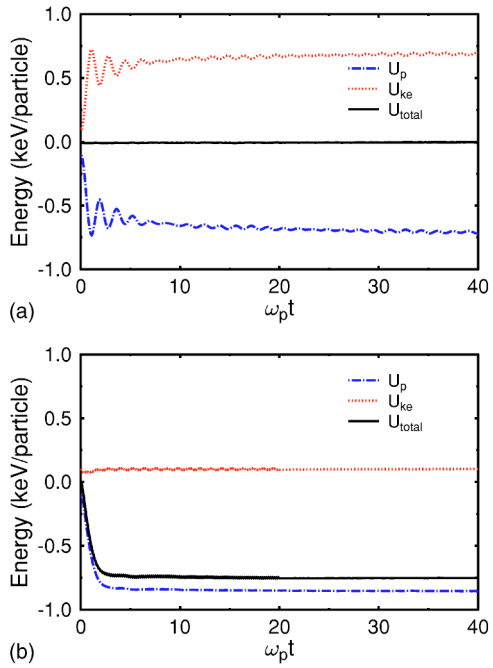

FIG. 2. (Color online). Relaxation of a cubic plasma consisting of 10 000 electrons and ions with (a) the total energy U_K+U_P conserved, and (b) with U_K held constant ($T_e=100$ eV) via a heat-bath correction to the equation of motion up to $\omega_p t=20$, and thereafter with the total energy conserved.

The end result, arrived at after a few plasma periods, is a configuration with well-defined boundaries, quasiuniform initial density, and minimum potential energy. The same potential energy U_P can also be reached by forcing total energy conservation (U_P+U_K), but only at the expense of increasing the electron temperature to some unpredictable value $\gg T_e=100$ eV, as demonstrated in Fig. 2(a).

It is important to note that unlike in conventional explicit PIC codes, the FSP model does not suffer from numerical heating associated with the grid instability (after all, there is no spatial grid here). The initial heating seen in Fig. 2(a) is *physical*, not numerical; the total energy (central line) is conserved—relative to the kinetic or potential energy values—to better than 1%, which is determined as much by the leap-frog integrator used as by the accuracy of the force summation.

The temperature-clamped system ($\omega_p t=0\rightarrow 20$) in Fig. 2(b) remains in thermal equilibrium when allowed to evolve in the absence of external fields. Note that in this case the potential energy ends up over eight times larger than the kinetic energy ($\omega_p t=20\rightarrow 40$), a situation normally associated with strongly coupled plasmas. For a charge-cloud plasma, however, the relevant parameter is $N_c=4\pi/3(\varepsilon/a)^3$, rather than $N_D=4\pi/3(\lambda_D/a)^3$, where a is the interparticle spacing. Although we still have to take some care over the choice of these parameters, the FSP model provides an effective means of modeling plasmas with finite, variable collisionality.

IV. LASER MODEL

Because wave propagation within the plasma is not yet supported by this model, the laser is incorporated by a ponderomotive source term, phase matched to the instantaneous critical density surface at the plasma edge. The appropriate amplitude and phase of the standing wave set up at the interface $x=x_c$ is determined by the solution of the Helmholtz equation for a normally incident, s-polarized plane wave on a step profile. Assuming zero absorption, this solution yields the following electric field:

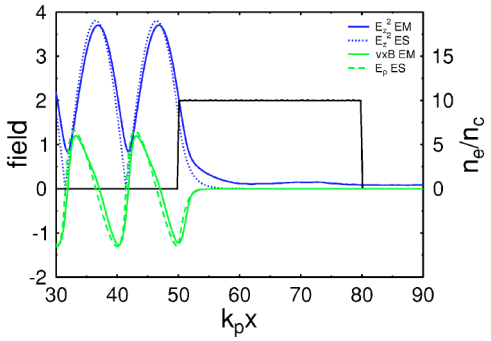

FIG. 3. (Color online). Ponderomotive laser model. A standing wave solution for the laser (amplitude $a_0=2.7$) is applied at the plasma–vacuum interface, giving rise to the intensity pattern E_z^2 (dotted curve) and a ponderomotive force (dashed curve). The equivalent fields computed from an electromagnetic PIC simulation are shown by the corresponding solid curves.

$$E_z=2E_L\cos\omega t\begin{cases}\sin(kx'+\phi), & x'<0,\\ \sin\phi\exp(-x'/l_s), & x'\geqslant 0,\end{cases} \tag{9}$$

where $\tan\phi=-kl_s$, $x'=x-x_c$, and $l_s=c/\omega_p$ is the collisionless skin depth. This field is assumed to maintain the above time dependence $\sim\cos\omega t$, which, after dropping the prime from the variable x (henceforth taken relative to the vacuum-plasma boundary) leads to the following expression for the x component of the $\mathbf{v}\times\mathbf{B}$ force on the electrons:

$$\begin{aligned} f_x^p &= v_z B_y \\ &= -E_z\frac{\partial E_z}{\partial x} \\ &= 2E_L^2\sin^2\omega t\begin{cases} k\sin[2(kx'+\phi)], & x'<0,\\ -\dfrac{2}{l_s}\sin^2\phi\exp(-2x'/l_s), & x'\geqslant 0.\end{cases}\end{aligned} \tag{10}$$

Note that unlike E_z, the ponderomotive force changes sign with x but *not* t; here it comprises an oscillating component at 2ω plus a dc component (the actual ponderomotive part), both of which always point in either the positive or negative x direction, as depicted in Fig. 3.

To make this laser model viable for three-dimensional, relativistic interactions, it needs two further modifications: a radial dependence to allow for a finite focal spot and a correction for large quiver amplitudes, $a_0=eE_L/m\omega c>1$. The expression used in the code therefore takes on the following form:

$$\mathbf{f}^p=-\nabla\gamma,$$

where

$$\gamma=(1+\Psi)^{1/2},$$

$$\Psi = 4a_0^2 X^2(x) R(r) T(t), \tag{11}$$

$$X(x) = \begin{cases} \sin\chi, & x<0, \\ \sin\phi \exp(-x/l_r), & x\geqslant 0, \end{cases}$$

$$R(r) = \begin{cases} \cos^2\left(\dfrac{\pi r}{4\sigma}\right), & r\leqslant 2\sigma, \\ 0, & r>2\sigma, \end{cases}$$

$$T(t) = \sin^2\left(\frac{\omega}{\omega_p} t\right). \tag{12}$$

The above expressions are written in terms of normalized variables, but for readability we have retained an explicit frequency ratio (ω/ω_p), so that the phase factors become: $\phi = -\tan^{-1}[(\omega/\omega_p) l_r]$, $\chi = (\omega/\omega_p) x + \phi$. The skin depth has been modified to account for enhanced penetration of the evanescent wave at relativistic pump strengths,[27] so that $l_r = \gamma_s^{1/2} l_s$, where $\gamma_s = (1 + 4a_0^2 n_c/n_u)^{1/2}$ corresponds to the field amplitude at the interface $x = x_c$, and n_u is an averaged upper shelf density (initially equal to n_0). Strictly speaking, this will render the equation for E_z invalid because it is also implied in γ_s; a more consistent option would be to use the formal solution for arbitrary a_0 and $n_0/n_c \gg 1$ given by Sudan.[28] Nevertheless, this first-order correction turns out to be quite good for the intensities considered here $(a_0 \leqslant 5)$. This is demonstrated in Fig. 3, which shows a comparison between the fields computed from the ponderomotive model [Eqs. (9) and (10)] and a fully electromagnetic, one-dimensional PIC code.[29]

The radial coordinate $r = (y^2 + z^2)^{1/2}$ is taken relative to the center of the focal spot, which is given a $\sin^2$—rather than a Gaussian—form in order to create a sharp radial cutoff at 2σ (σ is the half-width, half maximum). This is found to give a more physically reasonable modeling of profile deformation, avoiding penetration of the low-intensity wings in the overdense plasma, which would tend to occur for a Gaussian focal spot.

The normalized longitudinal and radial field components are finally given by

$$E_x^p = \frac{\partial\gamma}{\partial x} = \frac{a_0^2}{\gamma} R(r) T(t) \times \begin{cases} \dfrac{\omega}{\omega_p} \sin 2\chi, & x<0, \\ -\dfrac{2\sin^2\phi}{l_r} \exp(-2x/l_r), & x\geqslant 0, \end{cases}$$

$$E_y^p = \frac{\partial\gamma}{\partial y} = \frac{a_0^2}{\gamma} T(t) X^2(x) \begin{cases} -\dfrac{\pi y}{4\sigma r} \sin^2\theta, & r<2\sigma, \\ 0, & r\geqslant 2\sigma, \end{cases}$$

where $\theta = \pi r/4\sigma$.

V. HOLE-BORING TEST: PIC VERSUS TREE

This obviously simplistic model cannot hope to match the rich array of physical phenomena accessible through a full solution of Maxwell's equations. Nonetheless, when combined with a rudimentary density-tracking routine to monitor the position of the critical surface x_c and the mean density n_u just behind it, it does serve rather well in reproducing some of the main features of hot electron generation and ponderomotive ion dynamics.

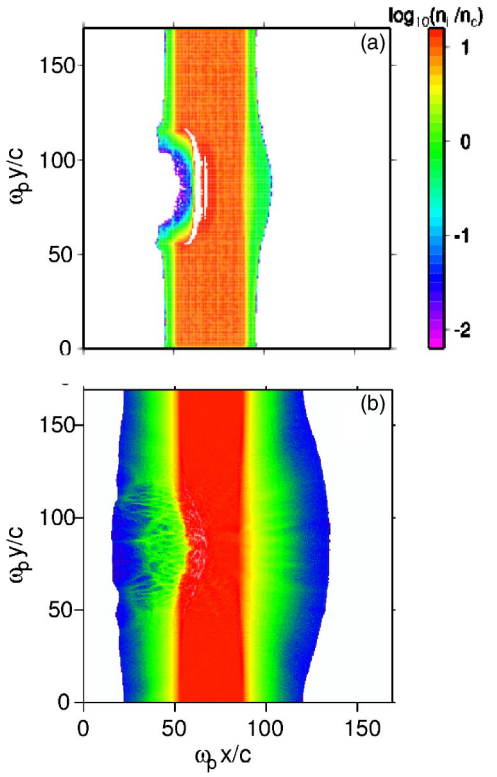

FIG. 4. (Color online). Ion density slice in the laser incidence plane at $700/\omega_p$ computed with (a) the tree code PEPC and (b) two-dimensional PIC simulation using OSIRIS. In (a) the density is averaged over the spot diameter in the z direction.

We illustrate this with a test problem in slab geometry, namely, collisionless shock formation through pressure imbalance—a hole-boring simulation. Balancing continuity and momentum at the critical surface (laser reflection point) gives the well-known formula for the recession velocity[30,31]

$$\frac{u_h}{c} = \left(\frac{Zm_e}{m_i}\frac{n_c}{n_e}\frac{2-\eta}{4} a_0^2 \cos\theta\right)^{1/2}, \tag{13}$$

where a_0 is the normalized laser amplitude or quiver velocity, η is the absorption fraction of laser energy coupled to the plasma, and θ is the angle of incidence.

A simulation to verify this behavior was set up using a plasma block with dimensions $60\times 170\times 60\ (c/\omega_p)^3$ and initial electron and ion temperatures of $T_e = 5$ keV and $T_i = 0$, respectively. The other simulation parameters were: $a_0 = 2.8$, $m_i/Zm_e = 1836$, $\eta = \theta = 0$, and $n_e/n_c = 10$. A total of 6.4×10^6 particles were used with effective size $\varepsilon = 2.5$ and average (ion) spacing $a = 0.5$, giving a smear factor $N_c = 4\pi/3(\varepsilon/a)^3 \simeq 600$ to ensure reasonably collisionless electron dynamics.

The shock and rear-side ion blowoff—features typical of this interaction scenario—can be readily identified in Fig. 4(a), which shows an ion density xy slice averaged over the spot diameter in the z direction after about 100 fs. For com-

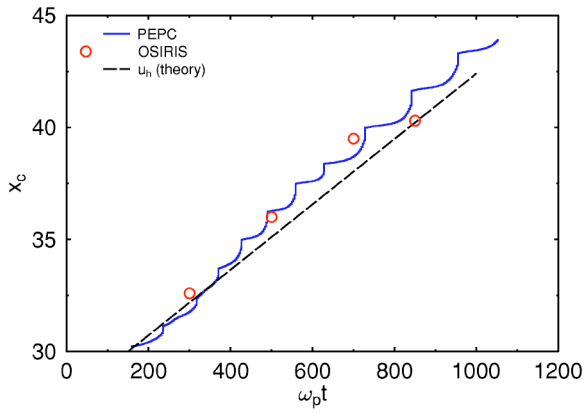

FIG. 5. (Color online). Hole-boring test: position of critical surface as a function of time in PEPC simulation (solid curve) and 2D PIC simulation performed with OSIRIS (circles) for a slab target irradiated at $I\lambda^2 = 10^{19}$ W cm^{-2} μm^2.

parison, a fully electromagnetic *two*-dimensional PIC simulation using the OSIRIS code[32] is shown in Fig. 4(b). In addition to the filamentary structures in the underdense region arising from propagation instabilities (which cannot be handled by the electrostatic model of Sec. II), the PIC simulation also exhibits faster ion blowoff than the tree code in this example. This is partly thanks to the better statistics (more simulation particles per unit area and therefore higher dynamic range in the ion density); but also an artifact of the periodic boundaries in the y direction combined with the large laser spot size, which allows fast electrons to recirculate across the simulation box. By contrast, the tree code employs open boundaries in all directions resulting here in less space charge in the vacuum regions.

Despite these differences in detail, the shock dynamics is very similar; Fig. 5 shows the position of the critical surface along the laser axis computed both from the tree code and from the PIC simulation. We see that both codes show good agreement with the theoretical value given by Eq. (13) of $u_h/c = 0.015$, giving us some confidence in the ponderomotive laser model described in Sec. IV.

VI. WIRE SIMULATIONS

The laser–wire simulations were set up with a quasineutral plasma with $n_i = n_e = (4,10) n_c (\omega/\omega_p = 0.5, 0.316)$ and initial electron and ion temperatures $T_e = 0.2-1$ keV and $T_i = 0$, respectively, configured in a cylinder with radius $R \simeq 1-4$ μm and height $H \simeq 10-16$ μm, as depicted in Fig. 1. Up to 3.2×10^6 simulation particles were used, with a mass ratio $m_i/m_e = 1836$ and ion charge state $Z = 1$. The laser wavelength is assumed to be 1 μm, the spot size is 0.5–1 μm full width at half maximum (FWHM), or $12-24\, c/\omega_p$ generally chosen so that the focal spot just fits within the wire diameter ($25\, c/\omega_p$). The pulse is turned on over five laser cycles and then kept at constant intensity for around 300 fs, or until the wire is burned through, at which point the standing wave ansatz described in Sec. IV is no longer reasonable. These parameters are still some way short of the experimental conditions, in which wires with diameters of 20 μm were irradiated by a 1 ps laser focused to 20 μm. The total laser energy converted into hot electrons is therefore 100–1000 times less in the simulations than in the experiment, so that we concentrate on identifying trends in the interaction behavior rather than attempting a rigorous quantitative comparison.

Scaling up the simulations is nontrivial because the statistics deteriorate rapidly; doubling the wire radius alone results in a quadrupling of the plasma volume $V = \pi R^2 H$, and therefore requires four times the number of particles to maintain the same particle macro charge Q_s (keeping n_e/n_c constant) and interparticle spacing (which is roughly $1\, c/\omega_p$ here). These parameters determine the maximum timestep permitted for numerical stability and hence the total simulation time required. The relatively low density and single ion species (protons) were also chosen for computational economy. A "minimal" simulation with $N_e + N_i = 1.44\times10^6$ particles took 50 h on 16 CPUs of the Jülich IBM p690+ Regatta. The largest simulation considered in the present work was actually the hole-boring test considered earlier: a $3\mu\times8\mu\times3\mu$ slab with 6.4×10^6 particles took over 100 h on 64 CPUs. For convenience we include a summary of the simulations referred to in Table I.

We begin our study by examining some general aspects of the laser–wire interaction for run C; the 2 μm ($50\, c/\omega_p$)-radius wire in the table. The large-scale electron and ion dynamics can be traced in Fig. 6, which shows a sequence of ion density slices in the x-z and x-y planes while the laser is incident. A number of features in (a) and (b) are immediately apparent: the strong bow-shock structure resulting from the ponderomotive push of the laser; the characteristic low-density ion blow-off back towards the laser; the electron return current along the wire axis from the tips towards the center (x-z) and from the wire surface back towards the focal spot (x-y); and ion layers starting to peel

TABLE I. Summary of simulation parameters: T_e is the initial electron temperature; N_c the particle "smear factor" controlling the collisionality; I_{19} the laser irradiance $I\lambda^2$ expressed in 10^{19} W cm^{-2} μm^2. U_L is the incident laser energy; U_e^a and U_i^a are the total energies absorbed by electrons and ions, respectively, at the end of the run; $U_i^{\max}$ is the maximum ion energy; U_i^{ave} the median ion energy (peak in spectrum).

RUN #	Dimensions $R\times H (c/\omega_p)$	n_e/n_c	T_e (keV)	N /10^6	N_c	Δt (ω_p^{-1})	I_{19}	τ_L (ω_p^{-1})	σ_L (c/ω_p)	U_L (mJ)	U_a^e (mJ)	U_a^i (mJ)	T_h (MeV)	$U_i^{\max}$ (MeV)	U_i^{ave} (MeV)
A	12×120	4	0.2	1.44	2	0.4	0.5	900	6	19	2.1	6.6	0.29	2.3	0.8
B	12×120	4	0.2	1.44	2	0.4	2.5	700	6	80	13	31	1.0	9	2.5
C	25×200	4	1	3.2	13	0.2	2.5	780	12	314	24	17	0.8	8	1.5
D	50×200	4	1	3.2	125	0.2	2.5	1450	12	600	400	120	0.7	6	1
E	25×200	10	0.2	1.44	2	0.4	2.5	900	6	21	1.5	4.8	0.33	2.5	1

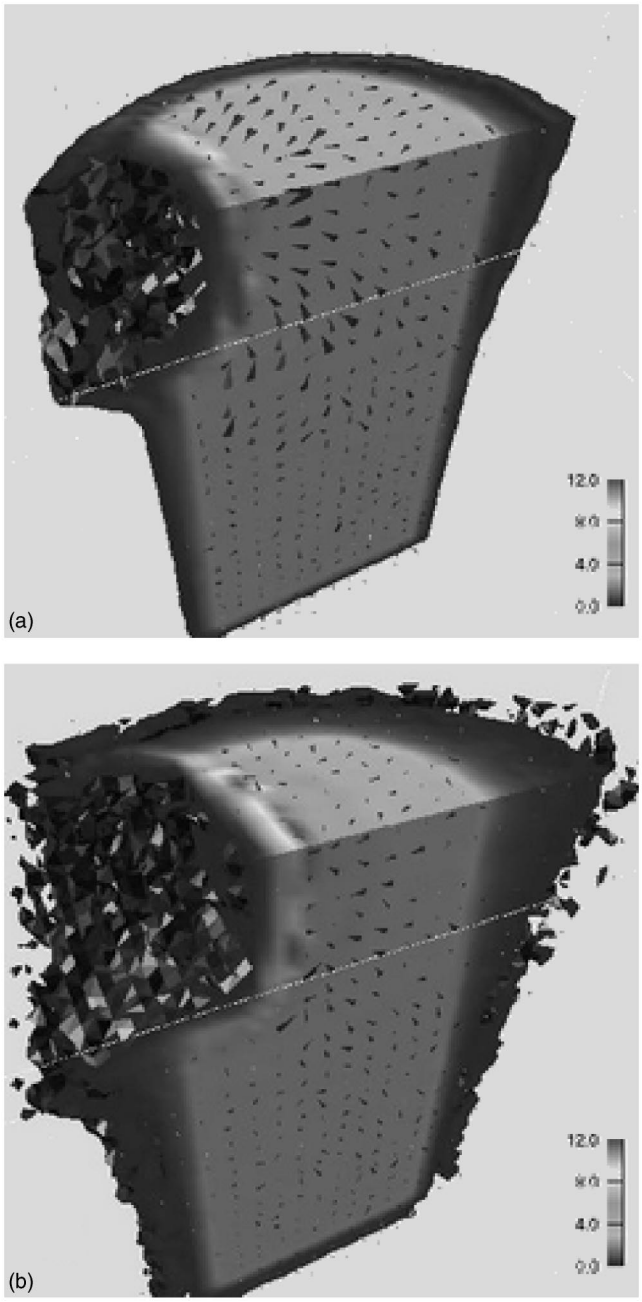

FIG. 6. Time-sequence of ion density isovolume $n_i/n_c \geq 0.1$ and electron return current j_e (arrows) for a 1/4 wire section sliced along the laser and wire axes, respectively—run *C*. Times shown are (a) $240/\omega_p$, (b) $800/\omega_p$, towards the end of the laser pulse. The laser is incident from the left.

off the rear side due to hot electrons (not shown) circulating behind and around the wire. This last effect is the familiar rear-surface acceleration mechanism, but in cylindrical geometry, ultimately leading to a disclike fast ion emission.

From Fig. 6 one might expect that rear-surface protons will dominate the emission spectrum here, yet this is only part of the picture. Inspection of the ion phase space (p_x-x) for Run *C* in Fig. 7(a) indicates that front-side ions are also accelerated significantly via the ponderomotively driven shock, some of which have already emerged from the rear surface (at $x=50$) as a beamlet in the forward direction. The onset of a double-disk structure is apparent in the p_z-x plot

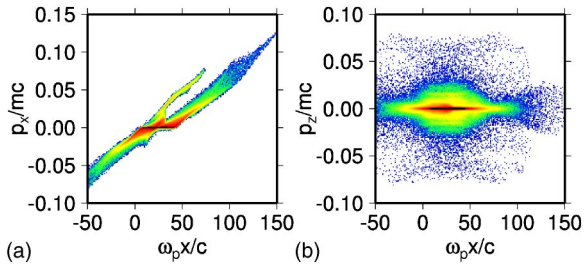

FIG. 7. (Color online). Ion phase space at the end of run *C* in the laser direction (x axis): (a) longitudinal and (b) vertical (p_z—along wire axis) momentum components. The wire is initially located between $x=0$ and $x=50$.

of Fig. 7(b): the rear-side ions ($x>100$) are beginning to fork at an angle of $5°-10°$ to the laser axis. At this point these ions have energies of >6 MeV, and are still being accelerated. Also evident from Fig. 7(b) are the significant blow-off components at $p_z \simeq \pm 0.05\, m_i c$ from the wire *tips*, reflecting the fact that the hot electrons have formed a large plume around the wire. Indeed, the electron phase space shows that this plume extends more-or-less symmetrically with a radius of $\sim 1200\, c/\omega_p$, or 50 times the initial wire radius. This corresponds to an effective simulation volume of almost $10^7\ \mu\text{m}^3$—a feat which would be difficult to match with a grid-based particle code.

The far-field structure of the ion emission in a more appropriate form for comparison with experimental measurements is shown in Fig. 8. Because the emission pattern in the y-z plane is not yet fully developed in Run *C* (this would require another 500 fs or so), we resort to a *predictive* diagnostic; namely, the angular momentum spread in the forward and backward directions. In other words, we compute the ion distribution $f(\alpha,\beta)$, where $\alpha=\tan^{-1}(p_y/p_x)$ and $\beta=\tan^{-1}(p_z/p_x)$. This is not quite the same thing as placing a virtual detector plate behind the wire, because the ions may still be undergoing acceleration—particularly in the y- and z-directions due to mutual repulsion—however, it does offer an early indication of the emission pattern. In Fig. 8(a) only the rear-side ions with energies >1 MeV are shown; the front-side ions, which initially form a radially symmetric beamlet with $\sim 10°$ spread, have been filtered out here. In (b) the ion blow-off back towards the laser is shown, which, as we see, also exhibits a stripelike emission pattern. These features are consistent with experimental data from the laser–

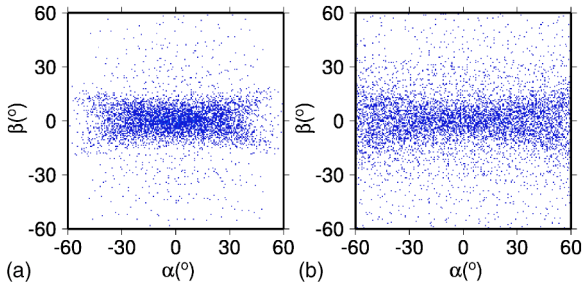

FIG. 8. (Color online). Angular MeV ion emission of (a) rear-side ions ($x>90\, c/\omega_p$) in the forward direction and (b) blowoff back towards the laser at the end of run *C*.

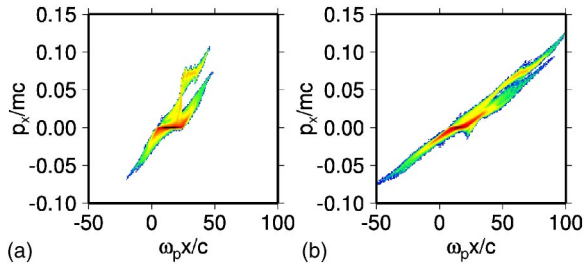

FIG. 9. (Color online). Ion phase space for the 1 μm wire in run B ($I\lambda^2 = 2.5\times10^{19}$ W cm^{-2} μm^2) at times (a) $t=600$ and (b) $t=1080$ showing enhanced front-side ion acceleration. The wire is initially located between $x=0$ and $x=25$.

wire experiments performed at RAL,[19,33] where emission was also observed over a large range of angles.

To get a feel for how these results scale with laser and target parameters, and to make a connection with the PIC simulations in Refs. 9–11, we have included a summary of the energy balance statistics in Table I.

Although this sample of the available parameter space is too small to draw definitive conclusions, some general trends are already apparent. First, the maximum ion energy $U_i^{\max}$ is clearly correlated to the laser intensity (or $I\lambda^2$) rather than the total power. The lower value for the 4 μm wire reflects the fact that most of the absorbed energy is either still carried by hot electrons or has gone into heating a larger bulk of wire material at this time. This is in contrast to the 1 μm wires, for which even after 700 ω_p^{-1} (200 fs), around 3$\times$ as much energy has been transferred to the ions than is carried by hot electrons.

The reason for this enhanced transfer efficiency is not clear at present. Normally, one would expect a smaller-radius wire to favor acceleration from the rear surface because the hot electrons have less material to pass through. However, runs A, B, and E have a far higher collisionality than C and D, implying a lower mean-free-path for the cold electrons. This in turn leads to inhibition of hot electron transport[34] and correspondingly more pronounced front-side ion acceleration. This can be clearly observed in the ion phase space of run B in Fig. 9, which shows the front-side ions emerging from the rear side with more than twice the energy than the rear-side ions.

A detailed analysis of the physics behind this effect will be presented elsewhere; for the time being, we compare the ion dynamics in the 1 μm radius wire with that observed in Fig. 6 for the 2 μm wire of run C. As before, we show a sequence of ion density isovolumes, but this time consisting

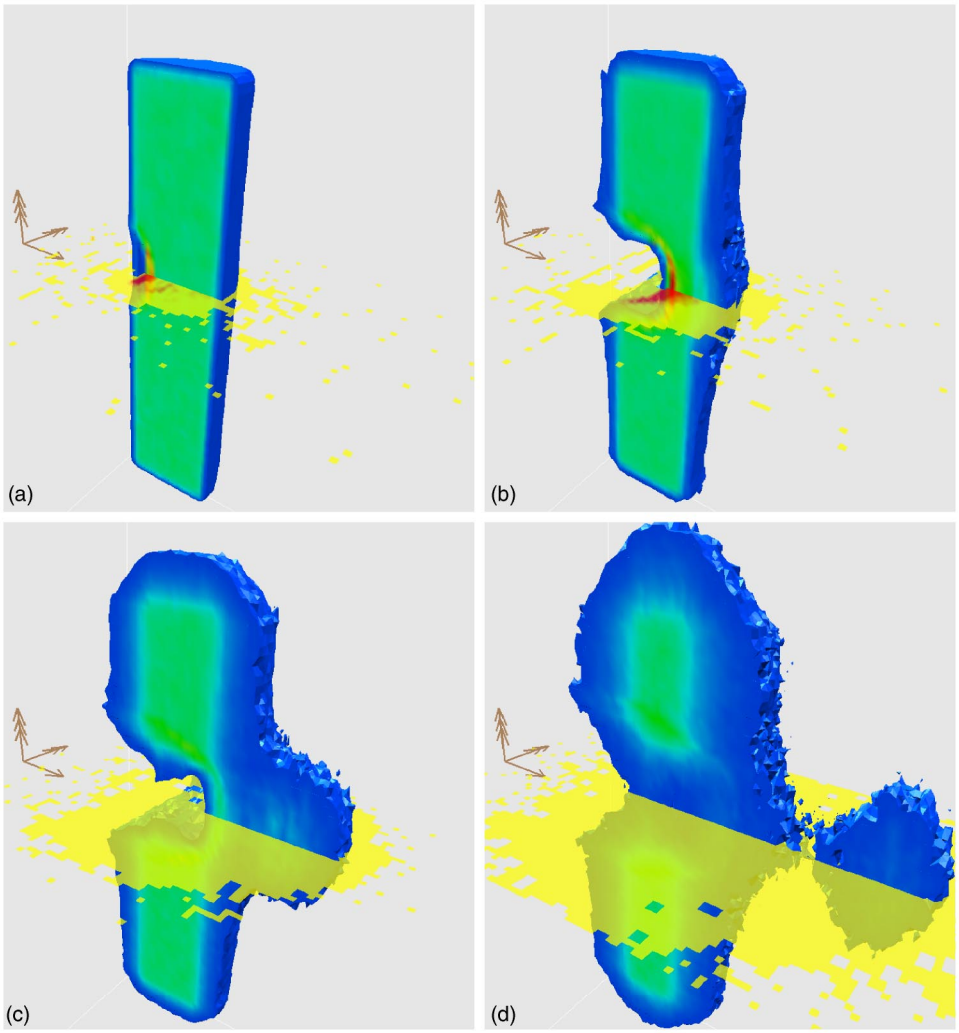

FIG. 10. (Color online). Time sequence of ion density isovolume $n_i/n_c \geqslant 0.25$ and electron temperature T_e slice in plane of laser incidence for a 1/2 wire section sliced along the wire z axis—run B. Times shown are (a) $200/\omega_p$, (b) $400/\omega_p$, (c) $600/\omega_p$, and (d) $800/\omega_p$.

of a 1/2 wire vertical slice—Fig. 10. Superimposed on these plots are slices of the instantaneous electron temperature, showing that while the laser is incident, the hottest electrons are actually confined to the shock region (a,b). At the same time, there is also a strong circulation of hot electrons around the wire.

The most striking feature of this simulation is that the entire mid section of the wire is pushed out by the laser; the beamlet visible in Fig. 10(d) has detached itself completely from the wire and continues to propagate away, spreading as it does so. This is reminiscent of three-dimensional PIC simulations of double-layer targets in which a proton beam was created from the low-Z coating on the rear side.[35] By contrast, the main thrust in this case comes unmistakably from the front side of the target, even though the beamlet comprises ions which originate from across the whole wire.

VII. DISCUSSION AND CONCLUSION

The simulations presented here demonstrate that high-intensity laser–wire interactions can be effectively modeled with a three-dimensional electrostatic tree code, despite simplifications to the absorption physics and the neglect of self-generated magnetic fields. The disclike ion emission pattern appears to originate initially from the cylindrically symmetric charge separation caused by hot electrons circulating around the wire. The vertical ''fine structure'' seen in Fig. 8 and in the recent experiments on VULCAN[33] appears to be caused by the large number of hot electrons arcing back onto the wire (which gets positively charged during laser irradiation), thus setting up a return current along the wire (z) axis from the tips to the focus. Ions exiting the wire surface are therefore pulled at a slight angle to the target normal, in the $\pm z$ direction for ions above and below the laser focal plane, respectively.

Return current effects also appear to play a role in the development of disklike emission in the small-radius wire simulations. In run B, for example, radial components develop in the aftermath of the burn-through phase, albeit at somewhat lower energies (0.5–1.5 MeV) in this case. Whether this effect persists as the wire radius and laser energy is scaled up will be addressed by future work.

ACKNOWLEDGMENTS

The tree-code simulations in this paper were performed on the IBM p690+ cluster, under Project No. JZAM04. Special thanks are due to the OSIRIS Consortium (UCLA/IST-Portugal/USC) for the use of their particle-in-cell code. The authors also acknowledge helpful comments and encouragement from A. E. Dangor and K. Krushelnick. P.G. wishes to thank M. Hammes (Wuppertal University) for assistance with the target preparation routines during the 2003 Guest-Student Program of the John-von-Neumann Institute for Computing.

[1] A. P. Fews, P. A. Norreys, F. N. Beg, A. R. Bell, A. E. Dangor, C. N. Danson, P. Lee, and S. J. Rose, Phys. Rev. Lett. **73**, 1801 (1994).
[2] F. N. Beg, A. Bell, A. E. Dangor *et al.*, Phys. Plasmas **4**, 447 (1997).
[3] T. Z. Esirkepov, V. S. Khoroshkov, A. V. Kuznetsov, and F. Pegoraro, Phys. Lett. A **299**, 240 (2002).
[4] G. Pretzler, A. Saemann, A. Pukhov *et al.*, Phys. Rev. E **58**, 1165 (1998).
[5] M. Roth, T. E. Cowan, M. H. Key *et al.*, Phys. Rev. Lett. **86**, 436 (2001).
[6] R. A. Snavely, M. H. Key, S. P. Hatchett *et al.*, Phys. Rev. Lett. **85**, 2945 (2000).
[7] A. J. Mackinnon, M. Borghesi, S. Hatchett *et al.*, Phys. Rev. Lett. **86**, 1769 (2001).
[8] M. Zepf, E. L. Clark, K. Krushelnick *et al.*, Phys. Plasmas **8**, 2323 (2001).
[9] S. C. Wilks, A. B. Langdon, T. E. Cowan *et al.*, Phys. Plasmas **8**, 542 (2001).
[10] A. Pukhov, Phys. Rev. Lett. **86**, 3562 (2001).
[11] H. Ruhl, S. V. Bulanov, T. E. Cowan, T. V. Liseikina, P. Nickles, F. Pegoraro, M. Roth, and W. Sandner, Plasma Phys. Rep. **27**, 363 (2001).
[12] E. L. Clark, K. Krushelnick, J. R. Davies *et al.*, Phys. Rev. Lett. **84**, 670 (2000).
[13] K. Krushelnick, E. L. Clark, M. Zepf *et al.*, Phys. Plasmas **7**, 2055 (2000).
[14] M. Zepf, E. L. Clark, F. N. Beg *et al.*, Phys. Rev. Lett. **90**, 064801 (2003).
[15] A. Maksimchuk, S. Gu, K. Flippo, D. Umstadter, and V. Y. Bychenkov, Phys. Rev. Lett. **84**, 4108 (2000).
[16] L. O. Silva, M. Marti, J. R. Davies *et al.*, Phys. Rev. Lett. **92**, 015002 (2004).
[17] E. L. Clark, Ph.D. thesis, University of London (2001).
[18] M. Roth, A. Blazevic, M. Geissel *et al.*, Phys. Rev. ST Accel. Beams **5**, 061301 (2002).
[19] F. N. Beg, M. S. Wei, A. E. Dangor *et al.*, Appl. Phys. Lett. **84**, 2766 (2004).
[20] P. Gibbon, Comput. Phys. Commun. (submitted), also available online at http://www.fz-juelich.de/zam/docs/printable/ib/ib-03/ib-2003-05.pdf
[21] S. Pfalzner and P. Gibbon, Phys. Rev. E **57**, 4698 (1998).
[22] A. B. Langdon and C. K. Birdsall, Phys. Fluids **13**, 2115 (1970).
[23] H. Okuda and C. K. Birdsall, Phys. Fluids **13**, 2123 (1970).
[24] T. Takizuka and H. Abe, J. Comp. Physiol. **25**, 205 (1977).
[25] S. Pfalzner and P. Gibbon, Comput. Phys. Commun. **79**, 24 (1994).
[26] S. Pfalzner and P. Gibbon, *Many Body Tree Methods in Physics* (Cambridge University Press, New York, 1996).
[27] P. Kaw and J. Dawson, Phys. Fluids **13**, 472 (1970).
[28] R. Sudan, Phys. Rev. Lett. **70**, 3075 (1993).
[29] P. Gibbon, A. Andreev, E. Lefebvre *et al.*, Phys. Plasmas **6**, 947 (1999).
[30] S. C. Wilks, W. L. Kruer, M. Tabak, and A. B. Langdon, Phys. Rev. Lett. **69**, 1383 (1992).
[31] K. G. Estabrook, E. J. Valeo, and W. L. Kruer, Phys. Fluids **18**, 1151 (1975).
[32] R. A. Fonesca, L. O. Silva, F. S. Tsung *et al.*, Lect. Notes Comput. Sci. **2331**, 342 (2002).
[33] F. N. Beg, M.-S. Wei, E. L. Clark *et al.*, Phys. Plasmas **11**, 2806 (2004).
[34] A. R. Bell, J. R. Davies, S. Guérin, and H. Ruhl, Plasma Phys. Controlled Fusion **39**, 653 (1997).
[35] T. Z. Esirkepov, S. V. Bulanov, K. Nishihara *et al.*, Phys. Rev. Lett. **89**, 175003 (2002).