Discontinuous Galerkin particle-in-cell simulation of longitudinal plasma wave damping and comparison to the Landau approximation and the exact solution of the dispersion relation

F. R. Foust,[1,a)] T. F. Bell,[1] M. Spasojevic,[1] and U. S. Inan[1,2]
[1]*Department of Electrical Engineering, Stanford University, Packard Building, 350 Serra Mall, Stanford, California 94305, USA*
[2]*Department of Electrical Engineering, Koç University, Rumelifeneri Yolu, 34450 Sariyer, Istanbul, Turkey*

(Received 6 May 2011; accepted 2 June 2011; published online 28 June 2011)

We present results showing the measured Landau damping rate using a high-order discontinuous Galerkin particle-in-cell (DG-PIC) [G. B. Jacobs and J. S. Hesthaven, J. Comput. Phys. **214**, 96 (2006)] method. We show that typical damping rates measured in particle-in-cell (PIC) simulations can differ significantly from the linearized Landau damping coefficient and propose a simple numerical method to solve the plasma dispersion function exactly for moderate to high damping rates. Simulation results show a high degree of agreement between the high-order PIC results and this calculated theoretical damping rate. © *2011 American Institute of Physics*. [doi:10.1063/1.3602123]

I. METHODS

Discontinuous Galerkin particle-in-cell (DG-PIC) has recently emerged[1] as a new approach to solving difficult problems in plasma physics. While, in principle, any Maxwell field solver may be used as a particle-in-cell (PIC) framework, in practice adoption of techniques other than finite-difference time domain (FDTD) has been slow. The nodal discontinuous Galerkin (DG) method,[2] however, is particularly interesting as a PIC framework due to its high order, high accuracy, and direct link between the solution at nodal DG points and polynomial interpolation at arbitrary points. We use the approach to demonstrate the range of validity of the Landau damping approximation; given the technique's high accuracy, we can readily show that the Landau damping rate agrees poorly with the actual observed damping rate under the conditions typically used for PIC simulations, that is, with moderate damping coefficients ranging from approximately -0.001 to -1.0. We therefore also present a simple numerical optimization approach to directly evaluating the hot plasma dispersion relation for longitudinal modes and show its essentially exact agreement with DG-PIC simulations under conditions where the Landau linearization fails.

A. The DG-PIC method

The technique and notation used here are nearly identical to that discussed in Ref. 1, with the exception of a minor sign error in their expression for the numerical flux function. We use the so-called hyperbolic cleaning approach[3,4] to enforce charge conservation, since the technique is trivially parallelized and does not require a global solve like Poisson projection techniques. We solve the augmented Maxwell's equations

$$\begin{aligned}\frac{\partial E}{\partial t}-\frac{1}{\epsilon}\nabla\times H+\chi\nabla\phi&=-\frac{1}{\epsilon}J,\\ \frac{\partial H}{\partial t}+\frac{1}{\mu}\nabla\times E&=0,\\ \frac{\partial\phi}{\partial t}+\chi\nabla\cdot E&=\chi\frac{\rho}{\epsilon}-\sigma_\phi\phi.\end{aligned}\tag{1}$$

The additional scalar field ϕ propagates out the divergence error at a speed χ. We also lightly damp this mode via the loss term σ_ϕ, effectively reducing the range of influence of a free charge. Provided this range is much larger than the effective Debye length, the loss term will not significantly affect the results and provides lower noise and greater predictability, particularly on periodic domains. The DG scheme begins by approximating the unknown solution $\boldsymbol{u}$ on an element e using basis functions $\phi_i(\boldsymbol{x})$

$$\boldsymbol{u}^e(\boldsymbol{x})\approx\sum_{j=1}^{N}\boldsymbol{u}(\boldsymbol{x}_i)\phi_i^e(\boldsymbol{x}).\tag{2}$$

We substitute this into the generic conservation law $\frac{\partial\boldsymbol{u}}{\partial t}+\nabla\cdot F(\boldsymbol{u})=0$, where F is a tensor, that is,

$$\nabla\cdot F(\boldsymbol{u})=\frac{\partial}{\partial x}F_x(\boldsymbol{u})+\frac{\partial}{\partial y}F_y(\boldsymbol{u})+\frac{\partial}{\partial z}F_z(\boldsymbol{u}).\tag{3}$$

Integrating by parts twice, we have the strong form of the DG scheme on an element e.

$$\begin{aligned}\int_V\left(\frac{\partial\boldsymbol{u}^e}{\partial t}+\nabla\cdot F^e(\boldsymbol{x})\right)\phi_i^e(\boldsymbol{x})\,d\boldsymbol{x}\\ =\oint_S\hat{\boldsymbol{n}}\cdot(F^e-F^*)\phi_i^e(\boldsymbol{x})\,d\boldsymbol{S}\end{aligned}\tag{4}$$

Finally, reversing the order of the summation and integration, we can carry out the integration directly for every pair of functions $(\phi_i(x),\phi_j(x))$ in the basis. This leads to the following matrix expression:

$$\frac{\partial u}{\partial t}+M^{-1}SF(u)=L(\hat{n}\cdot F(u)-\hat{n}\cdot F^*(u)),\tag{5}$$

where u is now understood to mean the vector of unknown weights for the basis. M is the mass matrix, whose entries M_{ij} are formed by the inner product of every pair $(\phi_i(x),\phi_j(x))$. S is the stiffness matrix, whose entries S_{ij} are formed by the evaluating the inner product of every pair

a)Electronic mail: ffoust@stanford.edu.

 © 2011 American Institute of Physics

Downloaded 04 Aug 2011 to 171.64.91.7. Redistribution subject to AIP license or copyright; see http://pop.aip.org/about/rights_and_permissions

$(\phi_i(x), \nabla\phi_j(x))$. The matrix L computes the surface integral of the flux contributions. In the nodal DG framework, the basis is chosen to be *interpolating* (specifically, the multidimensional Lagrange polynomials), so the scheme directly yields the nodal values at interpolation control points defined within each element. Details, in particular for Maxwell's equations, can be found in Ref. 2.

For the augmented Maxwell's equations listed above, the necessary flux $F^e - F^*$ can be computed by applying the Rankine-Hugoniot conditions across the normal and generalizing the resulting flux to include both full upwinding and central fluxes in one scheme. Details for the Maxwell system can again be found in Ref. 2. For the time being, we assume that the material parameters ϵ and μ may have jump discontinuities. The resulting fluxes are

$$\begin{aligned}
\hat{n}\cdot(F_E - F_E^*) &= \frac{1}{2\epsilon^-\{Z\}}(\beta\hat{n}(\hat{n}\cdot[E]) - \beta[E] \\
&\quad - Z^+\hat{n}\times[H]) + \frac{1}{2}\chi\hat{n}([\phi] - \hat{n}\cdot[E]), \\
\hat{n}\cdot(F_H - F_H^*) &= \frac{1}{2\mu^-\{Y\}}(\beta\hat{n}(\hat{n}\cdot[H]) - \beta[H] \\
&\quad + Y^+\hat{n}\times[E]), \\
\hat{n}\cdot(F_\phi - F_\phi^*) &= \frac{1}{2}\chi(\hat{n}\cdot[E] - [\phi]). \qquad (6)
\end{aligned}$$

The normal $\hat{n}$ is defined outward from the face of an element. The superscript " $-$ " denotes values interior to an element while " $+$ " denotes values exterior to the element. The notation $\{A\}$ means the average of the interior and exterior elements, that is, $(A^+ + A^-)/2$, and the notation $[A]$ denotes the jump across a face, $A^- - A^+$. Noting that if $Z^+ = Z^-$, that is, the materials are continuous across elements, the expression reduces to that in Ref. 1, with the exception of a neglected factor of $1/2$ and a minor sign error in their expression.

The parameter β controls the level of dissipation, from full upwinding ($\beta = 1.0$) to a purely central flux ($\beta = 0.0$). In the full upwinding case, we are implicitly assuming that the propagating characteristics of the *hyperbolic* part of Maxwell's equations are a good approximation to the unknown solution, an infinitesimal timestep in the future and in an infinitesimal neighborhood near the discontinuity between elements. Using these assumptions and also noting that the solution itself is discontinuous across the boundary, we can treat the problem as a simple initial value problem with discontinuous initial conditions. The Rankine-Hugoniot conditions then give the exact solution as a sum of characteristics propagating away from the interface at their characteristic velocities, which themselves can differ between the interior and exterior elements. Therefore, provided the assumption of hyperbolicity holds, discontinuities are handled explicitly and (in a limiting sense) exactly. In the case of the central flux ($\beta = 0.0$), we make no assumptions about propagating characteristics and simply assume that the intermediate solution is equal to the average of the exterior and interior fields. The material properties across the interface, however, are not a simple average but are, with appropriate adjustments for the differing units of E and H, the harmonic means of the characteristic impedance and admittance interior and exterior to a given element.

We recognize that in a nonlinear system such as a fully self-consistent plasma simulation, shocks may develop and it may be necessary to use a more dissipative flux or use limiters to ensure that ringing does not corrupt the solution. Since we are limiting ourselves to a study of the linear Landau damping rate in this paper, smooth initial conditions will yield a smooth solution; so, such fluxes are unnecessary and would simply add additional, unnecessary dissipation. In practice, many different choices of fluxes will lead to consistent schemes. The exact choice is not absolutely critical. For instance, we have found that our results are extremely insensitive to the choice of β. Aside from slightly more noise at high wavenumbers when $\beta = 0.0$, the solutions are effectively identical and the quantitative conclusions we draw below are unchanged, regardless of the choice of numerical flux.

The particles are tracked in 3 momentum dimensions and projected onto each element's control point using a generic shape function $S(r)$

$$J(x) = \sum v_i q_i S(|x - x_i|) \qquad (7)$$

$$\rho(x) = \sum q_i S(|x - x_i|) \qquad (8)$$

We accelerate the grid point lookup using a kd-tree or a simple spatial lookup table (a coarse grid). Which technique is more efficient depends a great deal on how uniform the grid is. Table lookups are faster for relatively uniform grids, where the kd-tree suffers from poor locality of reference. We have implemented this point search procedure for a fully three-dimensional grid; however, we note, for the 1D results in this paper, the spatial lookup table reduces to a simple hash table and the kd-tree reduces to a bisection search.

The particles in turn are moved by reconstructing the electric and magnetic fields at the point locations and evaluating the Lorentz force on the particles at those points. This amounts to evaluating a weighted sum over each function $\phi_i(x)$ in the given basis or equivalently (in the nodal DG framework) interpolation with polynomials through the fixed nodal control points. We accelerate this procedure by precomputing the polynomial coefficients on a reference element for a given polynomial order N. Since the particles are free to move between elements, we also require a lookup table for the particle containment search. We use either a kd-tree or a coarse grid, similar to the technique mentioned above, but storing each element's bounds. Again, the technique is implemented in 3D, but for the purposes of the 1D results in this paper, the search reduces to either a hash table lookup or a bisection search.

The system is parallelized using the Petsc MPI library[5–7] and stepped forward in time using either a split implicit-explicit IMEX Runge-Kutta scheme[3,8,9] or explicit Runge-Kutta. The system is typically stiff in the *field* solver for the type of problem discussed in this paper, so we choose that as the implicit component of the IMEX solve; consequently, each Runge-Kutta stage requires that a large, sparse system of equations be solved. For small problems, we use the Mumps LU solver.[10,11] For larger systems it is necessary to

Downloaded 04 Aug 2011 to 171.64.91.7. Redistribution subject to AIP license or copyright; see http://pop.aip.org/about/rights_and_permissions

use iterative methods. We find the BiCGSTAB (BiConjugate Gradient STABilized) iterative method[12] with the Euclid preconditioner[13,14] performs exceptionally well for the Maxwell DG system, acceptably converging in fewer than 10 iterations for moderately sized problems. The particles are integrated in the explicit part of the IMEX step; no global solve is required for the particle mover.

B. Damping rate calculations

Noting that the longitudinal electron mode cannot affect the distribution perpendicular to the direction of propagation, we use the purely one-component Maxwellian distribution function

$$f(v) = N\left(\frac{m}{2\pi k_B T}\right)^{\frac{1}{2}} \exp\left(-\frac{mv^2}{2k_B T}\right). \tag{9}$$

The linear Landau damping rate is derived by approximately solving the plasma dispersion function for longitudinal modes[15]

$$\epsilon(\omega,k) = 1 - \frac{q^2}{\epsilon_0 mk^2}\int_{-\infty}^{\infty}\frac{\partial f/\partial v}{v-\omega/k}dv = 0. \tag{10}$$

The Landau derivation begins by assuming that the frequency ω has a very small imaginary component ω_i representing loss. The singularity at $v = \omega/k$ is handled by integrating along a contour in complex v space. In typical derivations, the Landau contour is chosen to take a half-circle detour into the causal half-plane around the singularity, as shown in Fig. 1(left). Expanding the denominator in a series and further assuming that $k\lambda_D \ll 1$ yields the Landau damping rate

$$\omega_i = -\sqrt{\frac{\pi}{8}}\exp\left(-\frac{3}{2}\right)\frac{\omega_p}{(k\lambda_D)^3}\exp\left(-\frac{1}{2(k\lambda_D)^2}\right). \tag{11}$$

However, for large linear damping rates, the imaginary component of the frequency may be comparable to the real component, so it is necessary to choose a different integration contour as shown in Fig. 1(right). The singularity then contributes a factor of $2\pi i$ as opposed to πi in the small ω_i limit. Specifically, our new expression is (where v in the integration below is strictly real)

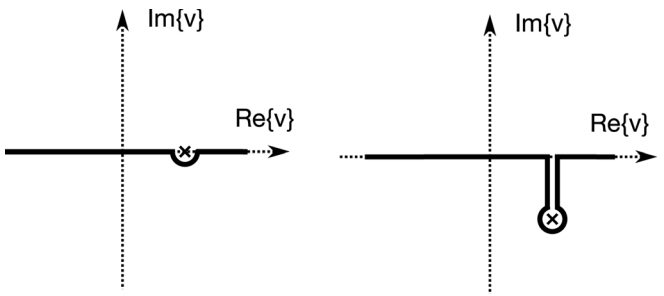

FIG. 1. Left: Landau integration contour in complex velocity space for very low loss rates. Right: integration contour for moderate to high loss rates.

$$\varepsilon(\omega,k) = 1 - \frac{q^2}{\varepsilon_0 mk^2}\left(\int_{-\infty}^{\infty}\frac{\partial f/\partial v}{v-\omega/k}dv + 2\pi i\frac{\partial f}{\partial v}|_{v=\frac{\omega}{k}}\right) = 0. \tag{12}$$

Of course, the two expressions agree in a limiting sense; however, as we will show in the next section, achieving good accuracy with the Landau linearization requires extremely small values of $k\lambda_D$, an unlikely regime for practical PIC simulations due to the extremely long integration times required.

We evaluate the integral at some ω and k by numerically integrating over some appropriate range of velocities (we choose 8 times the thermal velocity). The trapezoidal rule works well for this integration, but other approaches could also be used. Finally, the dispersion relation is solved as an optimization problem. In the half-plane $\omega_i < 0$, the dispersion function, as a function of ω, is generally smooth and well-behaved within the region of interest. We therefore solve the system by minimizing

$$g(\omega) = \log|\epsilon(\omega,k_0)| \tag{13}$$

using a standard direct search method for unconstrained optimization over (ω_r,ω_i). We have found that while the function is generally smooth and exhibits a clear minimum for the Maxwellian distribution, the minimum can exist in a long, narrow valley in (ω_r,ω_i) space when the loss rates are very small. Using the logarithm reduces the relative narrowness of the valley and this in turn helps reduce the overall number of iterations in the optimization routine. The true root should be in the vicinity of the Bohm-Gross and Landau frequencies (lossless longitudinal wave frequency ω_{r0} and damping rate ω_{i0}, respectively), so we choose the initial guess at $(\omega_{r0},\omega_{i0})$.

II. VALIDATION

The combination of high-order spatial accuracy with high-order time integration does not necessarily guarantee a quality PIC technique; however, prior work such as that of Jacobs *et al.*[1] have shown that even though neither energy nor momentum are explicitly conserved in the scheme, the scheme still has favorable numerical properties at least on par with, and sometimes exceeding, the performance of traditional FDTD-PIC. In particular, Jacobs *et al.* have shown that the DG-PIC technique shows an effective absence of particle self-force, low grid heating, and generally favorable numerical properties. To show the validity of the technique in the context of the Landau damping experiments in this paper, we first justify our choice of parameters in the next section by demonstrating the limit where DG-PIC fails.

A. Grid heating

PIC simulations often exhibit a slow rise in the total energy in the system, termed "grid heating." Grid heating represents a failure of the underlying fixed grid to represent high spatial frequencies in the particle population. Essentially, the error is caused by aliasing due to failure to sufficiently sample fluctuations on the scale of the Debye length,

Downloaded 04 Aug 2011 to 171.64.91.7. Redistribution subject to AIP license or copyright; see http://pop.aip.org/about/rights_and_permissions

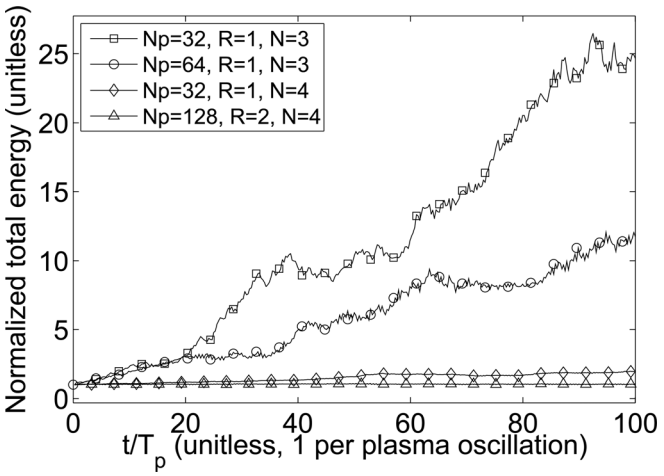

FIG. 2. Plot of the normalized total energy (TE = KE + PE) for a randomly-sampled Maxwellian electron plasma against an immobile, neutralizing background. Four cases are shown for various selections of the number of particles per element N_p, the particle radius as a function of the element size R, and the polynomial order N. These simulations are deliberately under-resolved in order to clearly show the effects of grid heating, with on the order of 20 Debye lengths per finite element.

manifesting as an exponential, aphysical energy rise in the system as a function of time.

In order to illustrate this failure, we choose parameters that are likely to give rise to very strong grid heating. We choose an electron temperature less than an order of magnitude smaller than in the next section, and discretize the same domain using only 25 elements instead of 100. Under these conditions, we have approximately 23 Debye lengths per element, that is, the simulation is significantly under-resolved and we would expect the artificial energy rise to be very large. Additionally, we initialize the system with relatively few particles, in the range of 32–128 per element, in order to ensure that grid heating manifests over a reasonable time scale. The results are shown in Fig. 2, where we plot the total energy (the sum of the particle kinetic energy and the potential energy in the electric field) as a function of time for a number of different parameters. N_p denotes the number of particles per element. R denotes the particle radius as a factor of the element size, that is, $R = 1$ means the particle cloud radius is the same size as one element. N denotes the polynomial order of accuracy. Our results are consistent with those in Ref. 1. Increasing the number of particles decreases grid heating slightly due to the higher smoothness of the original distribution. Increasing the particle radius has the same effect for the same reason. The most dramatic reductions in the magnitude of the artificial energy rise, however, had by increasing the polynomial order (and consequently, the number of unknowns per unit length) from $N = 3$ to $N = 4$. Similar results are obtained by increasing the number of elements from 25 to 50. If we increase all three parameters, the scheme exhibits practically zero energy rise over 100 plasma oscillations.

For our final Landau damping results below, we use a polynomial order $N = 4$, radius $R = 2$, number of particles in the range of $N_p = 1024$ to $N_p = 4096$, and 100 finite elements for a temperature more than an order of magnitude larger than the temperature used to generate the results in Fig. 2. Even at the lowest temperature, the worst case is 4 elements per Debye length, thus, based on these numerical tests, we do not expect grid heating to corrupt the solution in any significant way.

B. Conservation of momentum

Jacobs *et al.*[1] have already demonstrated conservation of momentum of the scheme on the microscopic scale, that is, the essential absence of particle self-force. In order to investigate global conservation of momentum, we initialize the system with a cold, drifting electron plasma in a neutralizing (constant opposing current, constant positive charge) background. A cold, drifting plasma is expected to be physically stable for all wavenumbers. Failure of either conservation of energy or conservation of momentum would manifest an instability in this simulation driven by inaccuracies as the particles drift over the finite grid. We choose the drift velocity such that a given particle will make approximately 16 transits of the domain over a period of 200 plasma oscillation periods. Again, we use a relatively low number of particles on a coarse grid, that is, the simulation is significantly under-resolved.

The results are shown in Fig. 3. Grid heating is present, manifesting as a very rapid initial phase of exponential

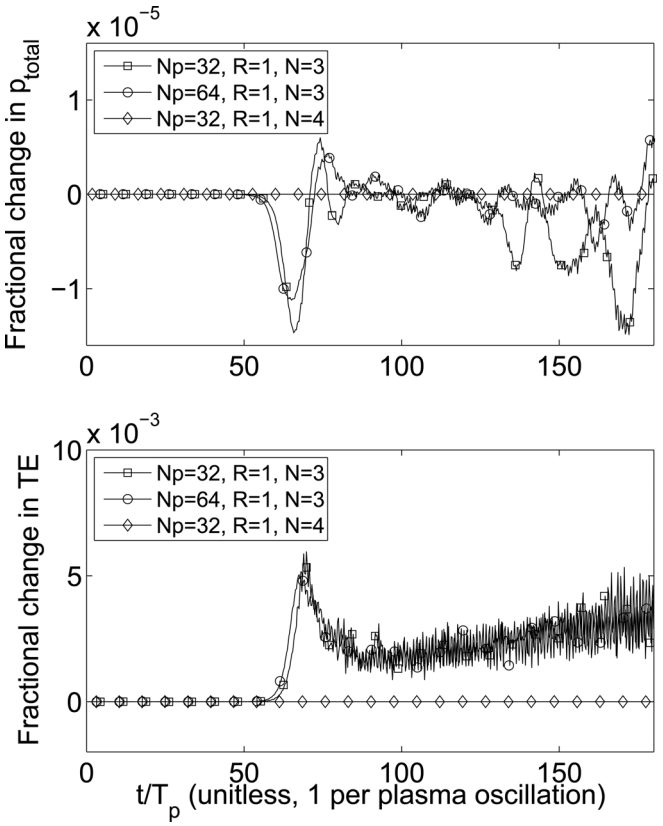

FIG. 3. Lower: fractional change in the total energy (TE = KE + PE) for a DG-PIC simulation of a cold beam drifting against a neutralizing background. Two unstable and one effectively stable case are shown. Upper: the fractional change in the time-averaged total momentum for the same simulation, plotted as a function of time. The time-averaged total momentum is calculated by windowing the instantaneous total momentum with a rectangular window of size equal to 4 plasma oscillation periods. The instantaneous momentum (not plotted) is dominated by small-scale fluctuations due to random particle motions.

Downloaded 04 Aug 2011 to 171.64.91.7. Redistribution subject to AIP license or copyright; see http://pop.aip.org/about/rights_and_permissions

growth in the total energy of the system. For reference, we plot this total energy in the lower panel. The first phase of growth (from approximately t = 0 to t = 60 plasma periods) is dominated by this exponential energy rise. Later times are dominated by individual particle fluctuations, a consequence of using too few particles. The instantaneous total momentum is dominated by small-scale fluctuations due to the small number of particles. Initially these fluctuations are on the order of the momentum of a single particle cloud. As the instability progresses, however, the beam is heated and these fluctuations in the instantaneous total momentum grow larger, just as we would expect based on the corresponding rise in the total kinetic energy. A global failure of conservation of momentum will manifest as a slow change in the *average* momentum of this beam as it transits through the domain due to viscous losses or artificial gains, that is, the beam will, en masse, either slow down due to artificial numerical friction or speed up for a similar reason. In order to clearly illustrate these long-term trends, we filter the instantaneous momentum using a rectangular window of width equal to 4 plasma oscillation periods. The results are shown in the upper panel, plotted as the fractional change in the average total momentum as a function of time. Despite momentum not being explicitly conserved in the formulation of the DG-PIC technique, it nonetheless performs adequately, showing only small fluctuations in the average momentum and no significant long-term loss or gain trends, even in the very unstable cases. Further, increasing the polynomial order from $N=3$ to $N=4$ stabilizes the simulation; no appreciable energy rise or change in the average total momentum is observed.

We now discuss the results of applying this technique to the simulation of Landau damping of the longitudinal electron mode.

III. RESULTS

We discretize a periodic domain using 100 finite elements and a 4th order polynomial basis. The fields, particle positions, and particle cloud projections are all periodized. To reduce the possibility of confounding numerical effects, we initialize the system with a very large number of particles, 1024 per element for low damping rates and 8192 for very high damping rates. The particles are loaded using a quiet start (ordered v-x space[16]). Each particle has the same charge-mass ratio but a different cloud weight, so particles on the long tail of the Maxwellian contribute proportionally less. The spacing in v space is chosen narrow enough to eliminate the multi-beam instability over the simulation time of interest. We simulate the system over a range of $\lambda_D k$, from approximately 0.25 (weakly damped, $\omega_i/\omega_p \approx -.002$) to 1 (strongly damped, $\omega_i/\omega_p \approx -1.0$).

The system is initialized with very small sinusoidal perturbation on the electron position, giving the initial (real) wavenumber k. The system is run long enough to observe significant damping. The frequency of oscillation is determined by peak detection, and the damping rate is determined by using an L1 norm log-linear fit to the peaks.

The results are summarized in Fig. 4. We show the Landau prediction and the damping rate as predicted by exact numerical evaluation of the plasma dispersion relation. The measured PIC simulation results are superimposed as discrete

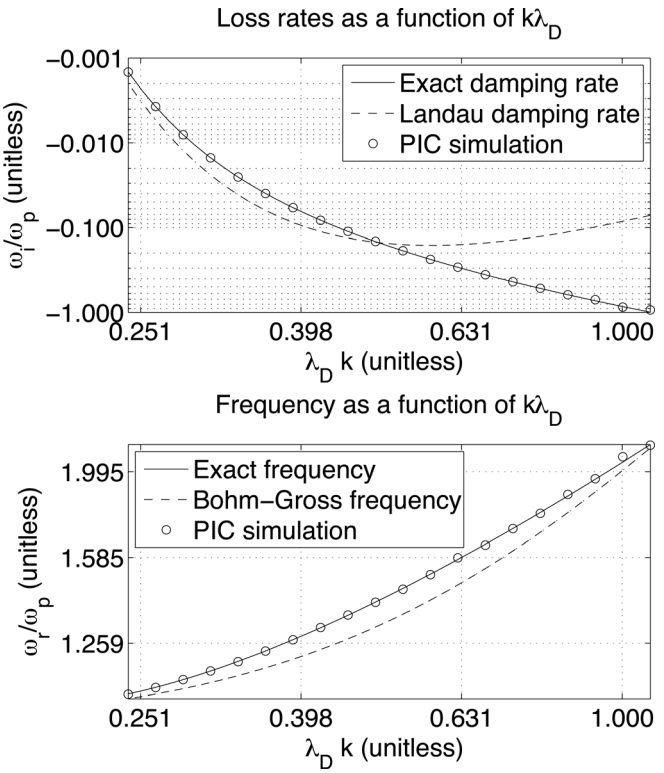

FIG. 4. Upper: log-log plot showing Landau damping rate, exact damping rate, and the measured damping rates from DG-PIC simulation for a Maxwellian distribution. Lower: same as above but for oscillation frequency.

circles. In all cases, the PIC simulation and the direct evaluation of the plasma dispersion relation show excellent agreement. Only at extremely low damping rates does the Landau approximation show good agreement with the simulation. The Bohm-Gross dispersion relation (yielding an estimate of the real component of ω) is similarly affected but less severely so.

IV. CONCLUSIONS

We have described a parallel implementation of a DG-PIC code and tested it in the context of Landau damping. We have provided a simple method to directly evaluate the plasma dispersion function and have shown excellent agreement with high-order DG-PIC simulations.

While very strongly damped modes may not be particularly amenable to observation in laboratory plasmas, they are commonplace in PIC simulations, where observing very weak damping rates would require prohibitively long integration times. We have shown that for moderate damping rates the exact plasma dispersion function is very easy to solve numerically, requiring no more than a few tens of iterations to converge to the solution and doing so yields damping rates that show very high agreement with high-order, high particle number PIC simulations.

In future work, we will extend these techniques to the simulation of transverse plasma modes.

ACKNOWLEDGMENTS

The authors were supported by the Office of Naval Research under Prime Award N000140710789 to the

Downloaded 04 Aug 2011 to 171.64.91.7. Redistribution subject to AIP license or copyright; see http://pop.aip.org/about/rights_and_permissions

University of Maryland with subcontract Z882802 to Stanford, and under Grant No. N00014-09-1-0034-P00003, and under National Science Foundation award 0840058.

[1]G. B. Jacobs and J. S. Hesthaven, J. Comput. Phys. **214**, 96 (2006).
[2]J. S. Hesthaven and T. Warburton, J. Comput. Phys. **181**, 186 (2002).
[3]G. B. Jacobs and J. S. Hesthaven, Comput. Phys. Commun. **180**, 1760 (2009).
[4]C.-D. Munz, P. Omnes, R. Schneider, E. Sonnendrücker, and U. Voß, J. Comput. Phys. **161**, 484 (2000).
[5]S. Balay, J. Brown, K. Buschelman, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang, see http://www.mcs.anl.gov/petsc for information on PETSc.
[6]S. Balay, J. Brown, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang, *PETSc Users Manual*, Tech. Rep. ANL-95/11 – Revision 3.1 (Argonne National Laboratory, Argonne IL 60439, 2010).
[7]S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, in *Modern Software Tools in Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen (Birkhäuser, Boston, MA, 1997), pp. 163–202.
[8]A. Kanevsky, M. H. Carpenter, D. Gottlieb, and J. S. Hesthaven, J. Comput. Phys. **225**, 1753 (2007).
[9]C. A. Kennedy and M. H. Carpenter, Appl. Numer. Math. **44**, 139 (2003).
[10]P. R. Amestoy, I. S. Duff, J. Koster, and J.-Y. L'Excellent, SIAM J. Matrix Anal. Appl. **23**, 15 (2001).
[11]P. R. Amestoy, A. Guermouche, J.-Y. L'Excellent, and S. Pralet, Parallel Comput. **32**, 136 (2006).
[12]H. A. van der Vorst, SIAM J. Sci. Stat. Comput. **13**, 631 (1992).
[13]D. Hysom and A. Pothen, SIAM J. Sci. Comput. **22**, 2194 (2000).
[14]D. Hysom and A. Pothen, in *Proceedings of the 1999 ACM/IEEE Conference on Supercomputing (CDROM)*, ISBN 1-58113-091-0 (ACM, New York, NY, 1999).
[15]L. Landau, J. Phys. USSR, **10**, 25 (1946) [JETP 16, 574]. Reproduced in *Collected Papers of L.D. Landau*, edited and with an introduction by D. ter Haar (Pergamon, Oxford, 1965), pp. 445–460;and in *Men of Physics: L.D. Landau*, edited by D. ter Haar (Pergamon, Oxford, 1965), Vol. 2.
[16]C. K. Birdsall and A. B. Langdon, *Plasma Physics Via Computer Simulation* (McGraw-Hill, New York, 1985).

Downloaded 04 Aug 2011 to 171.64.91.7. Redistribution subject to AIP license or copyright; see http://pop.aip.org/about/rights_and_permissions