HEAVY ION FUSION SCIENCE VIRTUAL NATIONAL LABORATORY

1ST QUARTER FY08 MILESTONE REPORT

REPORT INITIAL WORK ON DEVELOPING PLASMA MODELING CAPABILITY IN WARP FOR NDCX EXPERIMENTS

A. FRIEDMAN, R. H. COHEN, D. P. GROTE, J-L. VAY,
with the HIFS-VNL

December 13, 2007

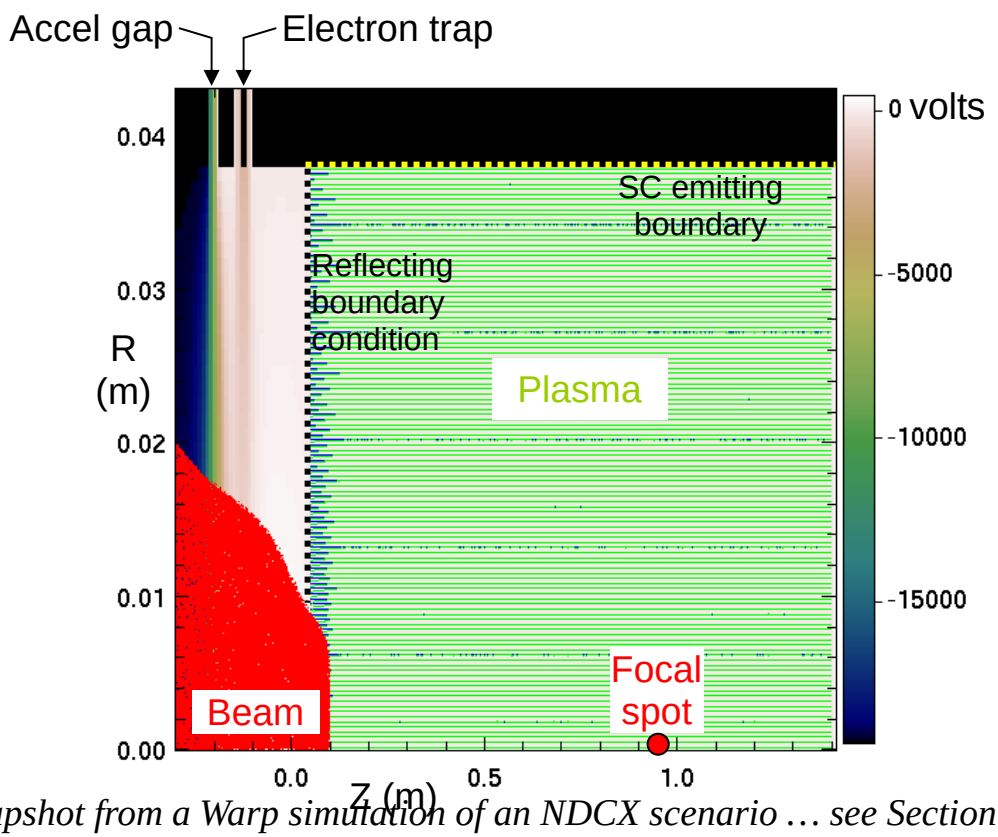

Snapshot from a Warp simulation of an NDCX scenario … see Section III D

Report on research carried out by the Heavy Ion Fusion Science Virtual National Laboratory in fulfillment of the FY08 First Quarter Milestone, December 2007:

"Report initial work on developing plasma modeling capability in WARP for NDCX experiments"

A. Friedman, R. H. Cohen, D. P. Grote, J-L. Vay, with the HIFS-VNL

December, 2007

SUMMARY

This milestone has been accomplished. The Heavy Ion Fusion Science Virtual National Laboratory (HIFS-VNL) has developed and implemented an initial beam-in-plasma implicit modeling capability in Warp; has carried out tests validating the behavior of the models employed; has compared the results of electrostatic and electromagnetic models when applied to beam expansion in an NDCX-I relevant regime; has compared Warp and LSP results on a problem relevant to NDCX-I; has modeled wave excitation by a rigid beam propagating through plasma; and has implemented and begun testing a more advanced implicit method that correctly captures electron drift motion even when time-steps too large to resolve the electron gyro-period are employed. The HIFS-VNL is well on its way toward having a state-of-the-art source-to-target simulation capability that will enable more effective support of ongoing experiments in the NDCX series and allow more confident planning for future ones.

I. INTRODUCTION

This milestone has been accomplished. The Heavy Ion Fusion Science Virtual National Laboratory (HIFS-VNL) has developed and implemented an initial beam-in-plasma modeling capability in Warp, including both an electrostatic implicit package and a novel (first of its kind) implicit "drift-Lorentz" package that more efficiently treats strongly magnetized electrons. These models will facilitate integrated simulations of present-day experiments, which are at sufficiently modest energies and currents that a fully electromagnetic description is not needed. The HIFS-VNL has also carried out a series of tests exploring the behavior of the models and confirming their proper action. Thus, the HIFS-VNL is well on its way toward having a state-of-the-art source-to-target simulation capability that will enable effective support of ongoing experiments in the NDCX series and allow confident planning for future ones.

The remainder of this report consists of two parts, the first describing model development and implementation, and the second describing simulation tests, both as relevant to this Milestone.

In Section II, we describe our new models. Subsection (A) describes the new direct-implicit model in Warp, and presents the motivation for this work. Subsection (B) describes the novel implicit drift-Lorentz model.

In Section III, we present the results of simulation tests. Subsection (A) compares the results of electrostatic and electromagnetic models when both are used to model beam expansion in NDCX-I relevant regimes, with and without neutralizing plasma, and with and without a magnetic field. This work was enabled by earlier development, especially the implementation of a conventional explicit fully electromagnetic package in Warp. Subsection (B) describes tests that confirm the stability of simulated plasma in a magnetic field, using the new implicit model with time steps larger than the plasma period. Subsection (C) describes our studies of numerically induced (spurious) heating of uniform plasma, a well-known phenomenon in PIC simulations that must be controlled for reliable results to be obtained. Subsection (D) describes simulations of combined longitudinal and transverse neutralized beam compression in plasma in a regime relevant to NDCX-I, including comparison of Warp and LSP code results. Subsection (E) describes two tests of the new implicit drift-Lorentz mover, one validating the implicit particle advance, the second confirming proper behavior when used to simulate (using a large timestep size that requires implicitness) the ion temperature gradient (ITG) mode in a magnetized plasma. Finally, subsection (F) describes a simulation of wave excitation by a beam passing through plasma, using Warp.

II. MODEL DEVELOPMENT

A. Direct-implicit model in Warp

The conventional explicit particle-in-cell (PIC) method suffers from a severe numerical instability when $\omega_{pe}\Delta t > 1$ and is thus usable only with small time steps. (Here, ω_{pe} is the electron plasma frequency and Δt the timestep size.) An implicit method, such as the direct-implicit algorithm[i], relaxes this constraint and is usable even when time steps that do not resolve plasma oscillations are employed. A version of this algorithm has been implemented in Warp. This implicit advance is stable with $\omega_{pe}\Delta t >> 1$, and has been tested successfully with $\omega_{pe}\Delta t > 100$. Such a capability becomes highly valuable as the plasma density increases and ω_{pe} becomes large. This stability comes at the price of a somewhat more difficult system to solve; other potential problems, such as energy non-conservation, must be avoided by careful choice of algorithmic details and simulation parameters.

The implicit advance contains a circular dependency – the future particle positions (that is, the positions at the advanced time level at the end of the time step) depend on the future fields, and these in turn depend on the future particle positions. The direct-implicit method self-consistently captures this dependency by linearizing the effect of the future field on the particles and incorporating it into the field solution. Thus, the future field acts on itself through the particles. (An effective susceptibility term is incorporated into the Poisson equation.) This formulation results in a slightly more complicated field solution compared with the explicit scheme, since the implicitness introduces non-constant coefficients into the matrix equation; the particle advance is changed only slightly. The linearization must include terms corresponding to the effect of the magnetic field (in our electrostatic case, just the applied magnetic field need be included). In particular, the method implemented in Warp uses the "simplified differencing" scheme to obtain the linearized charge density, and uses two-cell averaging of the implicit susceptibility – both schemes are described in the text by Birsdall and Langdon[ii]. The field-solver has been implemented in 2-D slab, 2-D axisymmetric, and most recently full 3-D geometries.

B. Implicit version of "drift-Lorentz" mover

For efficient simulation of magnetized plasmas as needed for the neutralized drift-compression of beams, it is desirable to be able to move particles (in particular, electrons) with timesteps Δt exceeding both the plasma period (inverse of the plasma frequency ω_p) and cyclotron period. Hence we have developed an implicit version of the electrostatic Drift-Lorentz particle mover. (Note that this element of our code development was supported by LLNL Laboratory-Directed Research and Development (LDRD) funds.) The Drift-Lorentz mover[iii] is a blend of the well-known "Boris" algorithm (which advances particles using the full Lorentz force) with drift kinetics (which advances particle guiding centers, as appropriate for a strong magnetic field). This novel "blended" mover can be used to advance particles in arbitrary-strength magnetic fields. The algorithm recovers correct particle drifts and a proper gyration radius for large $\omega_c\Delta t$, and full particle dynamics in the limit of small $\omega_c\Delta t$, where ω_c is the cyclotron frequency. It was originally developed as an explicit scheme and so was limited to low density ($\omega_p\Delta t$). An initial step to remove this restriction was described previously[iv]: a two-timestep predictor-corrector scheme incorporating a portion of the polarization drift in the Poisson equation. That procedure introduced implicitness into the dynamics perpendicular to the magnetic field. Results for a test problem, an ion-temperature gradient (ITG) instability, were presented in that reference. It was shown that known results from gyrokinetic simulation (a technique applicable to strongly magnetized plasmas) could be recovered, but that at least two corrector steps were required for large-timestep stability; also, the simulation geometry employed in those earlier studies precluded the development of numerical instabilities with wave vectors purely along the magnetic field.

We have now formulated and implemented a fully implicit treatment that closely parallels the direct-implicit algorithm described above. Specifically, particles are initially advanced (with the Drift-Lorentz mover) neglecting the contributions of the potential at the advanced timestep; the resulting positions are used to construct a linearized correction to the charge density from the (unknown) advanced potential, which is included in the "left-hand side" of the Poisson equation, and the resultant fields are then used to correct the particle positions. A stability analysis predicts that, for the ITG test problem, the previous mover (with polarization in the field solve) should be unstable (to period $2\Delta t$ oscillations) with a single corrector, as observed, while the new implicit scheme is predicted to be stable.

III. SIMULATION TESTS

A. Beam expansion in vacuum and plasma, comparing results from electrostatic and electromagnetic models.

Using parameters in the general regime of NDCX-I, we simulated a beam "slab" impinging on a plasma in 2-D (x,z) planar geometry, using both the explicit electrostatic (ES) model and the explicit electromagnetic model (EM) in Warp. For the EM case, subcycling of the electromagnetic field advance (that is, use of a smaller field-advance timestep size) was employed, so that the Courant stability condition for the field advance would be satisfied while the particle-advance timestep was kept the same for the two models (EM and ES). In these runs, the beam was initially overlaid with an artificial "K^-" beam moving to the left, as an artifice to ensure that the initial system was charge neutral; this was purely for convenience. As seen in Fig. 1, the results of the EM and ES models are indistinguishable from each other, confirming our expectation that, for experiments in the modest energy and current regime of NDCX-I, self-magnetic and inductive effects have an insignificant influence on beam behavior.

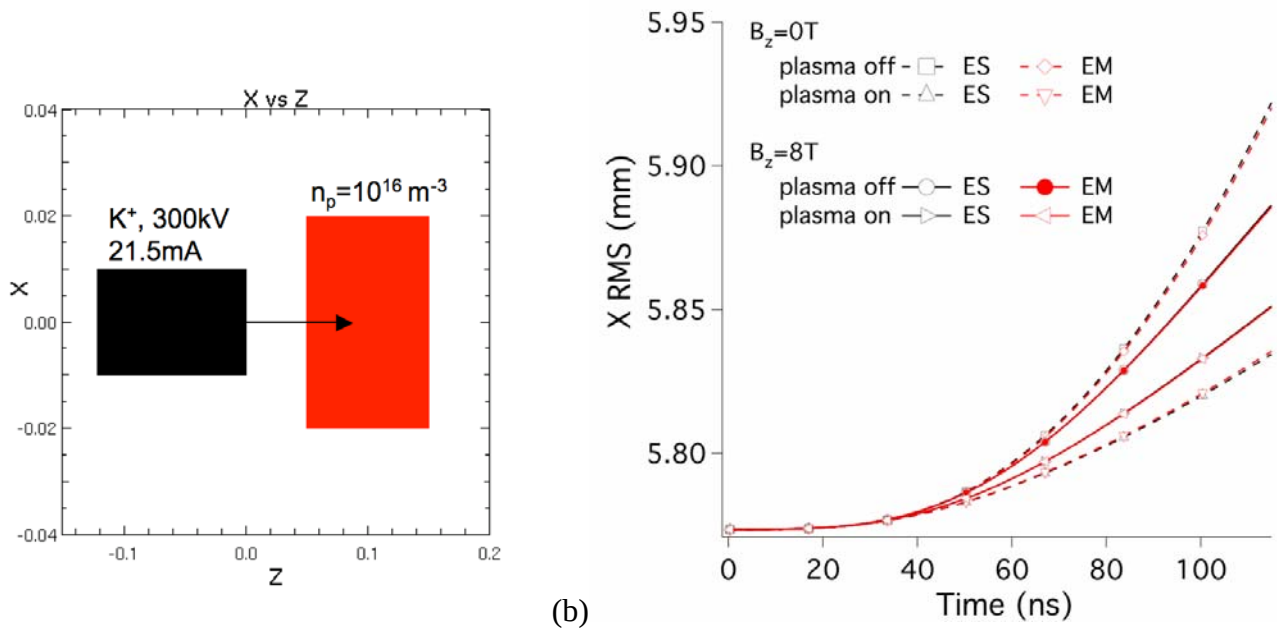

Figure 1. (a) initial configuration of Warp 2-D "slab" *(x,z)* simulations of the expansion of a beam in vacuum and in a plasma. The beam and plasma shapes are respectively 2 cm wide × 12 cm long, and 5 cm wide × 10 cm long rectangles. Both the beam and plasma are initialized with uniform density profiles at respectively 3.5×10^{14} m^{-3} and 1×10^{16} m^{-3}, providing a large reservoir of electrons for beam space-charge neutralization; (b) history of the instantaneous transverse RMS size of the entire beam using an electrostatic field solver ("ES") or an electromagnetic field solver ("EM") with or without plasma ("plasma off/on"), and with or without an external uniform longitudinal magnetic field B_z = 8T.

It is important to note that, even for a "large" system (doubling the original size of the simulation domain along each dimension), the boundary condition plays a critical role in determining the forces that act on the beam. Figure 2 shows this dependence, for the larger system. Note that the duration of these runs corresponded to many light-transit times across the simulation domain.

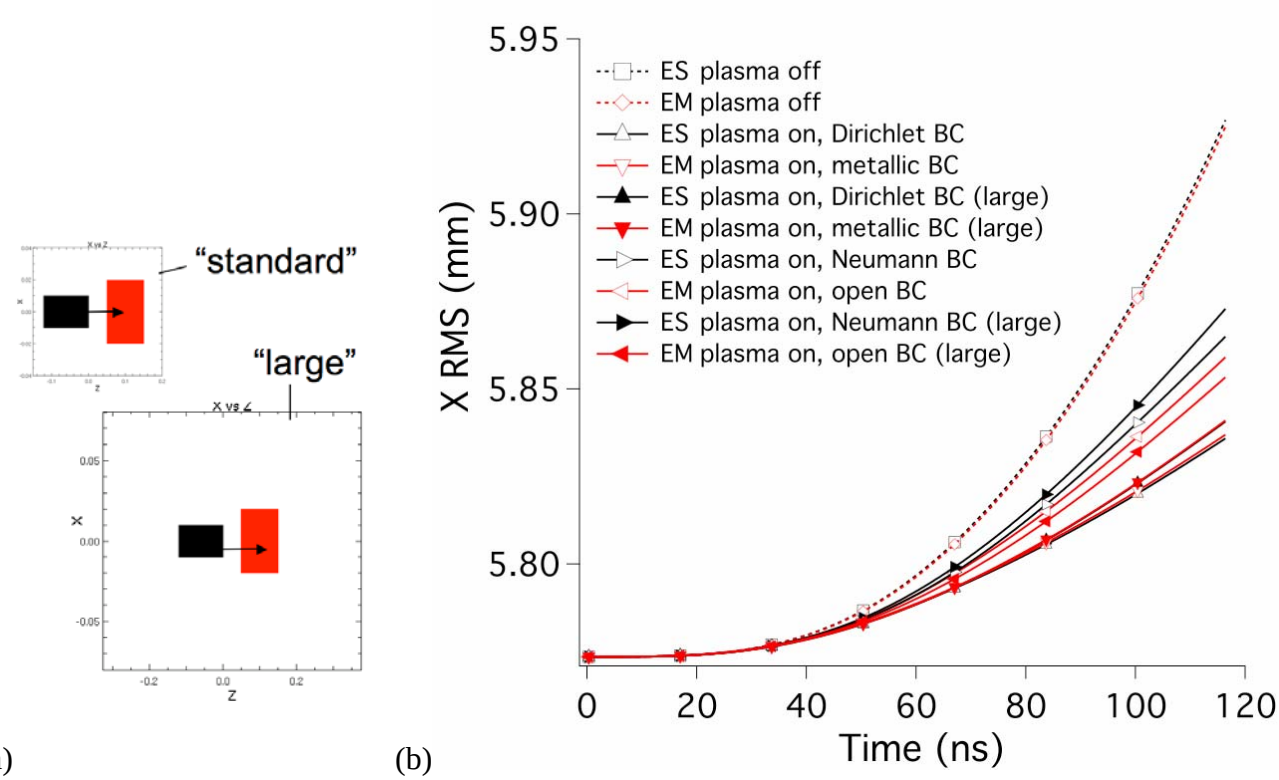

(a) (b)

Figure 2. (a) initial configuration of Warp 2-D "slab" *(x,z)* simulations of the expansion of a beam in vacuum and in a plasma for two sizes of the simulation domain: "standard" (8 cm × 35 cm) and "large" (16 cm × 70 cm); (b) history of the instantaneous transverse RMS size of the entire beam using an electrostatic field solver ("ES") or an electromagnetic field solver ("EM"), using different boundary conditions and different sizes of the simulation domain.

B. Numerical stability when simulating plasma in magnetic field with $\omega_{pe}\Delta t > 1$

A basic test of the direct-implicit model for a magnetized plasma was performed, demonstrating stability for a uniform plasma in a uniform magnetic field with $\omega_{pe}\Delta t > 1$. Parameters were chosen so that the electron cyclotron motion was (marginally) resolved, while electron plasma oscillations were not. The problem was set up in x-z geometry, with the magnetic field in the y-z plane (y being the symmetry direction), a configuration adopted from the ITG test problem for the blended mover noted in Sec. IIB. The specific parameter choices were: $T_i = T_e = 9.7$ KeV, $B_y = 100$ G, $B_z = 0.01\ B_y = 1$ G, density $n = 2.1\times10^{16}$ m^{-3}, and $\Delta t = 5.2\times10^{-9}$ s. The cell size was $\Delta x = \Delta z = 7$ mm, and the simulation domain was a square of 16 × 16 cells. The ions were protons, and the electrons had an artificial mass equal to the proton mass/100. For these parameters, $\omega_{pe}\Delta t = 10$ and $\omega_{ce}\Delta t = 0.5$, while $\omega_{pi}\Delta t = 1$ and $\omega_{ci}\Delta t = 0.005$. The electron gyroradius/cell size was 14.1. The simulation was performed with 65,536 each of electrons and ions. The problem was run with both implicit and explicit (standard Boris) timestepping (same timestep). As predicted, the implicit scheme was observed to be stable, with the normalized RMS electrostatic potential $(e\Phi/T_e)_{rms}$ remaining small, ~ 10^{-5}, over a run of duration 5000 ω_{pe}^{-1}; see Fig. 3. In contrast, an explicit run with the same parameters showed severe numerical instability, with $(e\Phi/T_e)_{rms}$ rising from 0.01 to 3 in seven timesteps, by which time over half of the electrons were lost from the system (by the route of being displaced over many cells in a single timestep, sufficient to wind up more than a grid-cell width outside the simulation volume).

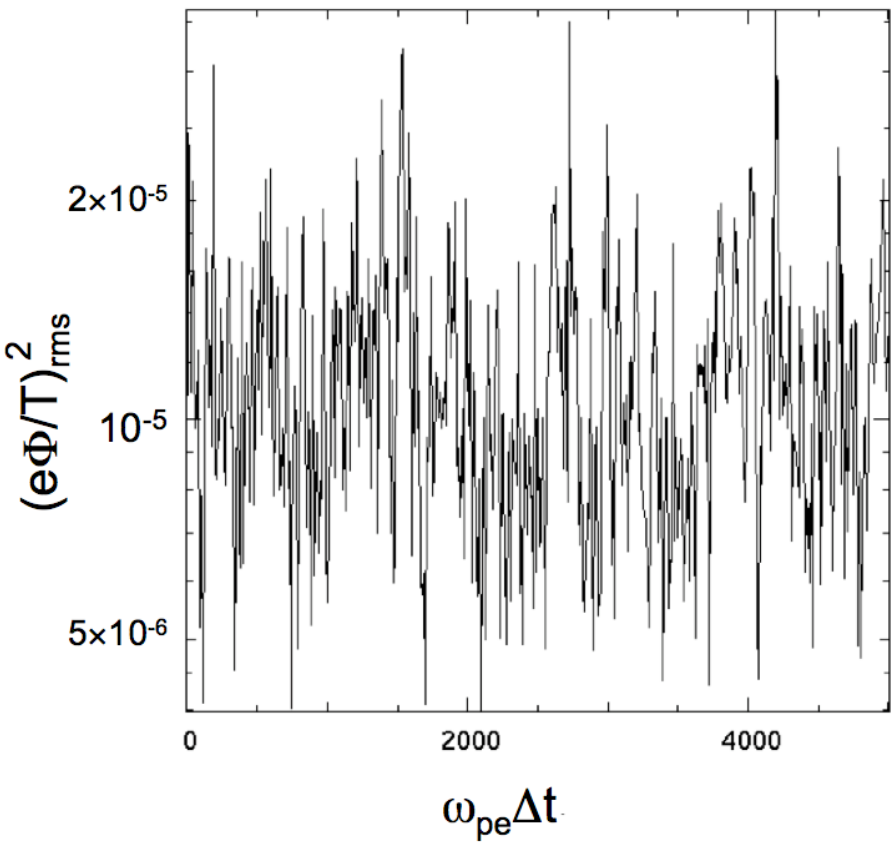

Figure 3. RMS normalized fluctuation amplitude versus time for implicit simulation of uniform plasma; explicit simulation with same timestep develops large potential fluctuations in a few timesteps, sufficient to cause substantial particle loss. Note the logarithmic scale of the ordinate.

C. Numerical heating of uniform plasma

Other basic tests of Warp's plasma simulation capabilities were carried out and comparisons were made with LSP results. The primary purpose was quantification of the numerically induced heating obtained using various simulation algorithms. The results were much as expected. Explicit simulations were carried out for a uniform plasma with doubly periodic boundaries in two dimensions. For the momentum-conserving scheme, with small enough grid cell size to resolve Debye lengths, $\Delta x = 0.5\lambda_{De}$, the simulations are "stable," showing little heating. With a larger grid cell size, $\Delta x = 5\lambda_{De}$, numerical heating is significant, increasing the temperature by a factor of over 100 after 1000 plasma periods. With the energy-conserving scheme, the simulations showed little heating with the larger grid cell size, for a sufficiently short run. These results agree well with the corresponding LSP simulations.

When the Warp simulations were run for longer time using the energy-conserving scheme, it was seen that, after 1000 or so plasma periods, the temperature suddenly grew rapidly, becoming larger than it had been with the momentum-conserving scheme. This late-onset rapid heating is believed to be due to violation of the particle Courant condition, resulting from the initial modest temperature increase. That slow heating causes $v_{th}\Delta t/\Delta x$ to exceed unity eventually, violating the Courant condition and leading to instability. With the step size decreased by a factor of 16, the Courant condition was not violated during the simulations, and no late onset heating was seen. The LSP simulations were not carried out long enough in time to see this effect.

D. Longitudinal and transverse neutralized compression

A primary motivation for development of plasma simulation models in Warp is to enable simulations of neutralized longitudinal compression and focus using the same tools that are used to simulate the un-neutralized beam in the injector and accelerator. We have begun carrying out simulations of beam longitudinal compression and transverse focusing in plasma, concentrating on understanding what models are needed in Warp and their domains of applicability, and on reproducing results of LSP simulations[v,vi]. The system simulated is an idealized model of some of the neutralized drift-compression and final-focus experiments on NDCX-1. A converging beam is injected at the left edge of the simulation region. The beam passes through a "tilt core" (accelerating induction gap) that initiates the longitudinal compression by applying a head-to-tail velocity variation. The beam then enters the neutralizing plasma. The convergence of the beam and the applied velocity tilt are adjusted so that the transverse focus coincides with the point of peak axial compression. The parameters are typical of those studied on NDCX-1: a singly charged potassium beam with energy 300 kV, current 18 mA, initial radius 2 cm, initial convergence 32.2 mrad, and temperature 0.2 eV. The focal length and compression length are approximately 0.9 m. The plasma is a singly ionized carbon plasma, with densities in the range from 10^7 to 10^{10} cm^{-3}, and temperatures from 0 to 50 eV.

In all cases, approximately nine plasma simulation particles of each species were initialized per cell, totaling 46170 particles for each species. Enough beam particles are injected so that there are initially 100 particles per cell, and the total for the full beam is 157688 particles. The transverse grid cell size was 1 mm, and the longitudinal size 1 cm, giving a grid that is 42x170 grid cells in r and z, respectively. The plasma Debye length was 0.18 mm, unresolved by the grid, and so the energy-conserving mover was used. The time step size was 0.0333 ns, giving $\omega_{pe}\Delta t = 0.188$.

A goal of these simulations is to understand how the plasma parameters and resulting neutralization affect the beam compression. However, the simulations have shown some sensitivity of the focal spot size to the transition region, where the beam enters the plasma. For example, any electrons that flow upstream will cause the neutralization to occur earlier than expected, leading to an earlier focus. In order to remove these sensitivities and produce a reliable comparison, a special boundary condition at the plasma entrance was applied in some cases. In Warp, a reflecting boundary is used, reflecting any plasma particles that would otherwise move upstream. The benchmark LSP simulations[vi] were carried out in two steps, to achieve the same effect. The first step used no plasma, and saved the beam data as the beam passed through the plane of the nominal plasma entrance. The second step was with plasma, starting at the plasma edge and recreating the beam from the saved data.

For a case with compression factor of 67 and spot size 1.3 mm, the beam density is 7.37×10^7 cm^{-3} upon injection and increases to 1.17×10^{12} cm^{-3} at focus. The base case examined has a plasma density of 10^{10} and an initial temperature of 0 eV (cold plasma). At this density, the focal spot is nearly optimal since there is nearly complete neutralization of the beam space charge. Good agreement with LSP results was obtained for both the longitudinal compression, a factor of 67, and the focal spot size, 1.3 mm. Figure 4 shows images ("snapshots") from the simulation, as the beam enters the plasma and just past the time of peak compression. Figure 5 shows Warp and LSP results for the spot size as a function of time at the approximate location of peak compression, $z = 0.95$ m. Note that, in this simulation, the special boundary condition at the plasma edge was not needed.

Figure 5 also includes results from a series of simulations done with lower plasma densities. These cases were run to examine the effects of imperfect neutralization. When the plasma density is less than the beam density, nearly all of the electrons are pulled into the beam. In these simulations (with both codes), space-charge-limited emission of electrons is enabled at the pipe wall in order to better model the ability of an externally-connected plasma to balance the beam's space charge. Qualitative

agreement with the LSP results was obtained, showing an insensitivity of the longitudinal compression ratio to the plasma density and an increased focal-spot size with reduced neutralization.The degree of agreement between the codes is governed by details of the model of the transition region where the beam enters the plasma, and by the boundary conditions. As discussed in the following paragraph, the techniques used in these runs to simulate an idealized plasma edge affected the beam dynamics.

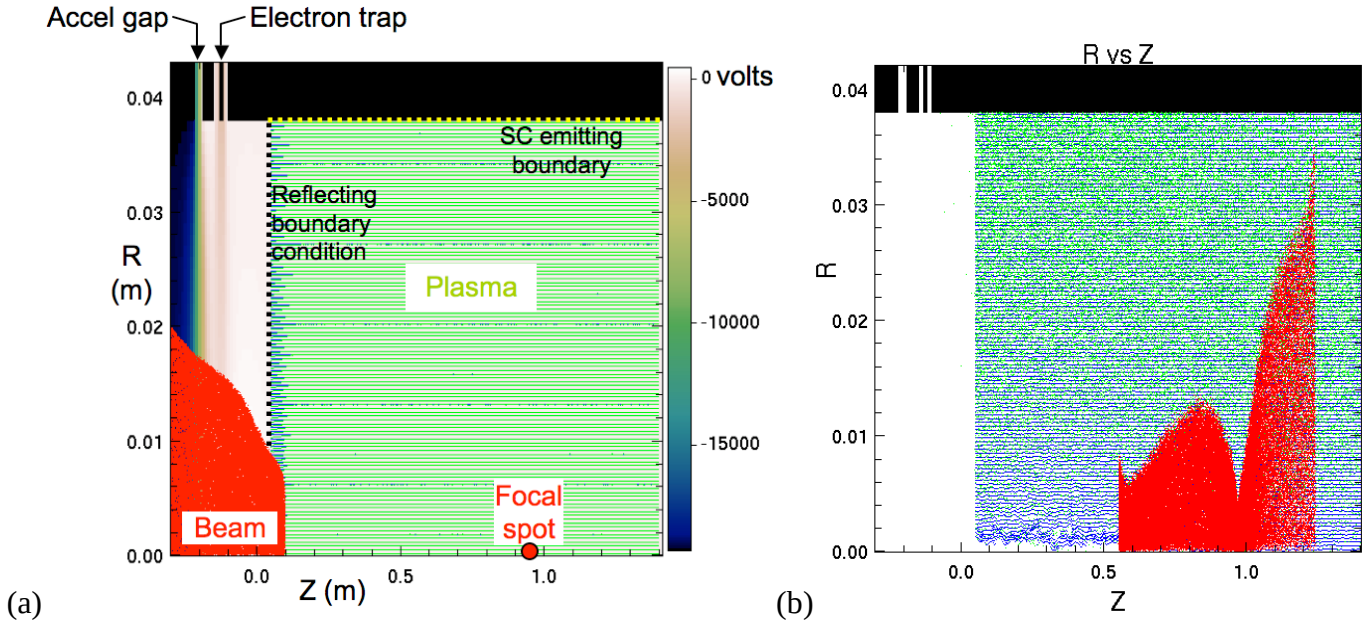

Figure 4. Snapshots of the beam and plasma (a) as the beam enters the plasma and (b) just past the time of peak compression and focus. The red dots are the beam, the green dots plasma electrons, and the blue dots plasma carbon ions. In (a) the background colors show the value of the electrostatic potential. Note the large negative value at the left edge – it is from the negative swing of the tilt waveform. In (b), most of the current is in the small region near z = 1 meter.

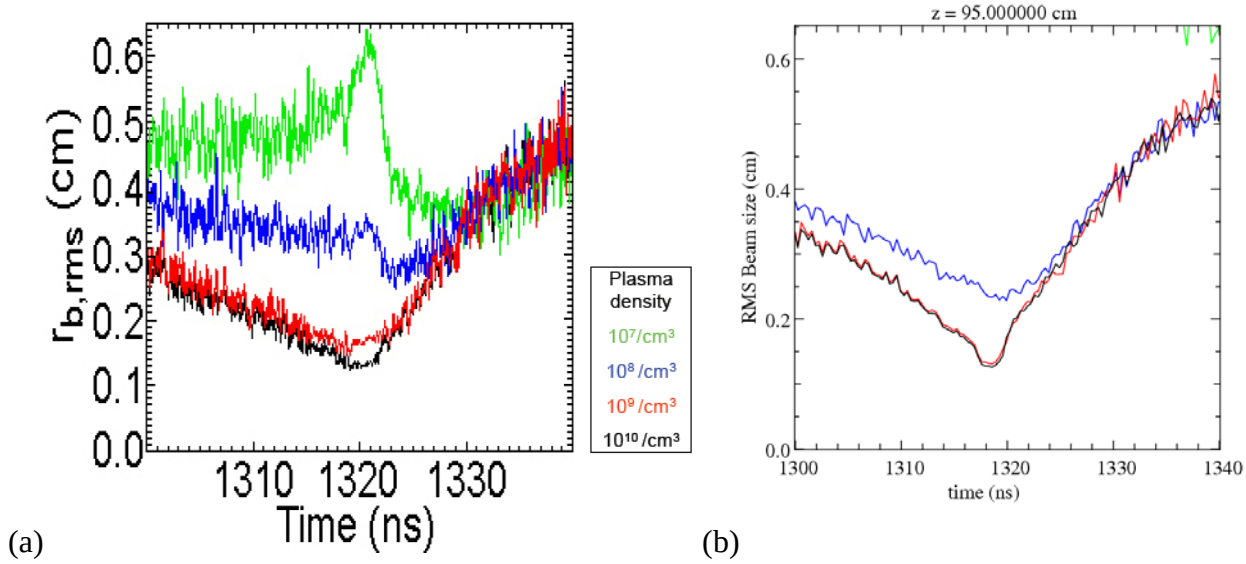

Figure 5. RMS beam radius as a function of time, at the axial location of peak longitudinal compression, z = 0.95 m, for various plasma densities: (a) LSP results; (b) Warp results. Note that in (b) the focal spot radius for the pathological 10^7 cm^{-3} density case (green) was much larger than was observed in LSP and is mostly off the plot. The graphic in (a) was taken from Ref. [6].

In the LSP simulations, a two-stage method was used, employing two separate runs to model the system. In the first stage, in which there was no plasma, the beam was injected and a tilt was applied by the gap. The beam data was saved as it passed the location of the plasma entrance. In the second stage, there was a plasma, and the beam was regenerated from the saved data at the plasma entrance. The Warp simulations described above used a single run to model the system, including the beam injection, application of the tilt, and the plasma. In pursuit of closer agrement between the codes, the same two-run method was implemented in Warp. Much closer agreement with the LSP results was indeed obtained. The results for the base case of 10^{10} plasma density are shown in Fig. 6. The difference between the one- and two-stage cases highlights the sensitivity of the beam behavior to the plasma boundary - but note that in this case, the spot size is not affected.

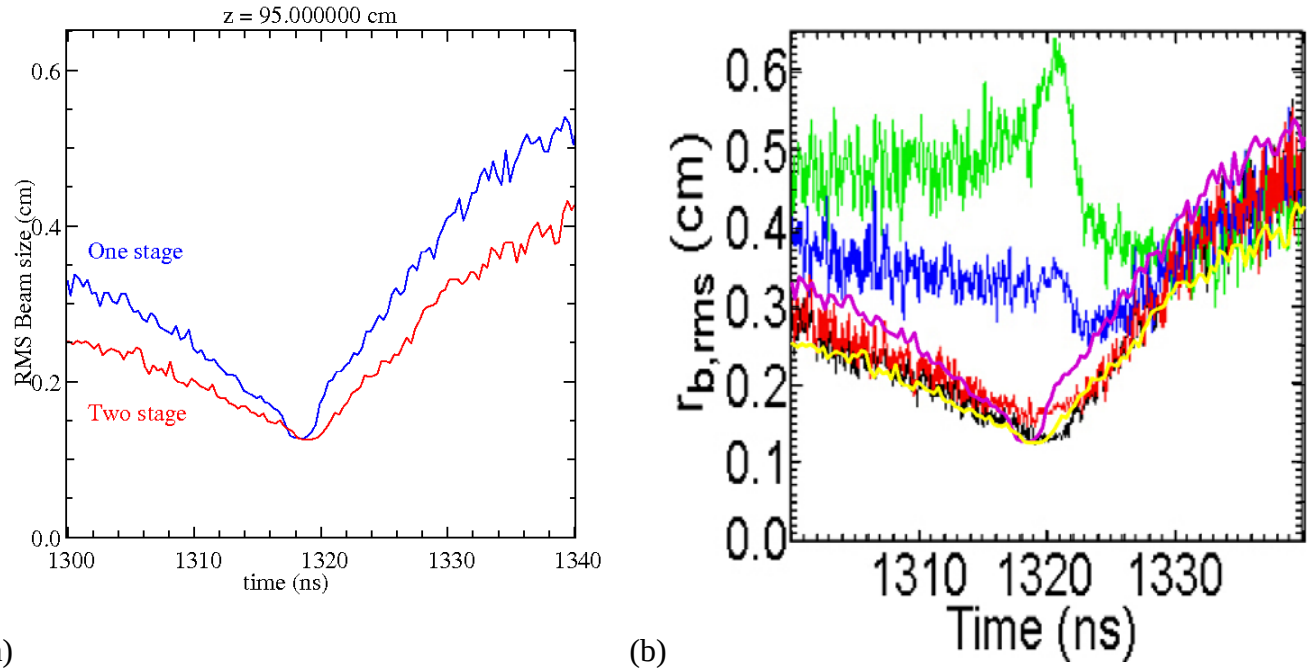

(a) (b)

Figure 6. RMS beam radius vs. time, at the axial location of peak compression, z=0.95 m, for different models: (a) the blue curve is the one-stage Warp simulation, and the red the two-stage Warp simulation; (b) Warp and LSP results overlaid; the yellow line shows the Warp two-stage result, to compare against the black line (from LSP). The one-stage Warp result is shown in magenta. Note that the Warp curves are less noisy because of down-sampling - the data was taken only every 10 steps, versus every step in LSP.

The Warp simulations described above were all carried out using the explicit model. The simulation of the base case, 10^{10} per cm^3 density and 0 eV temperature, was also carried out using the implicit model, to test that capability. The implicitness allowed a 15-fold larger timestep, giving $\omega_{pe}\Delta t = 2.82$. These runs required less than an hour of computer time each, in contrast with (typically) a day for the explicit runs. Good agreement was found for the peak longitudinal compression, but the beam reaches its minimum spot size earlier. A comparison of the time-histories of the beam radius (at the location of peak longitudinal compression for the explicit runs) is shown in Fig. 7. A larger radius is seen for the implicit case at the point of peak axial compression since the beam is beginning to rebound after its earlier focus. It is not clear at this time what is causing the more rapid focusing. The figure also shows the peak current (maximized over time), and the RMS beam radius at the time of that peak current, versus axial location, for both explicit and implicit runs. Note that the LSP run was done with the implicit electromagnetic field solver, but with an explicit advance for the particles.

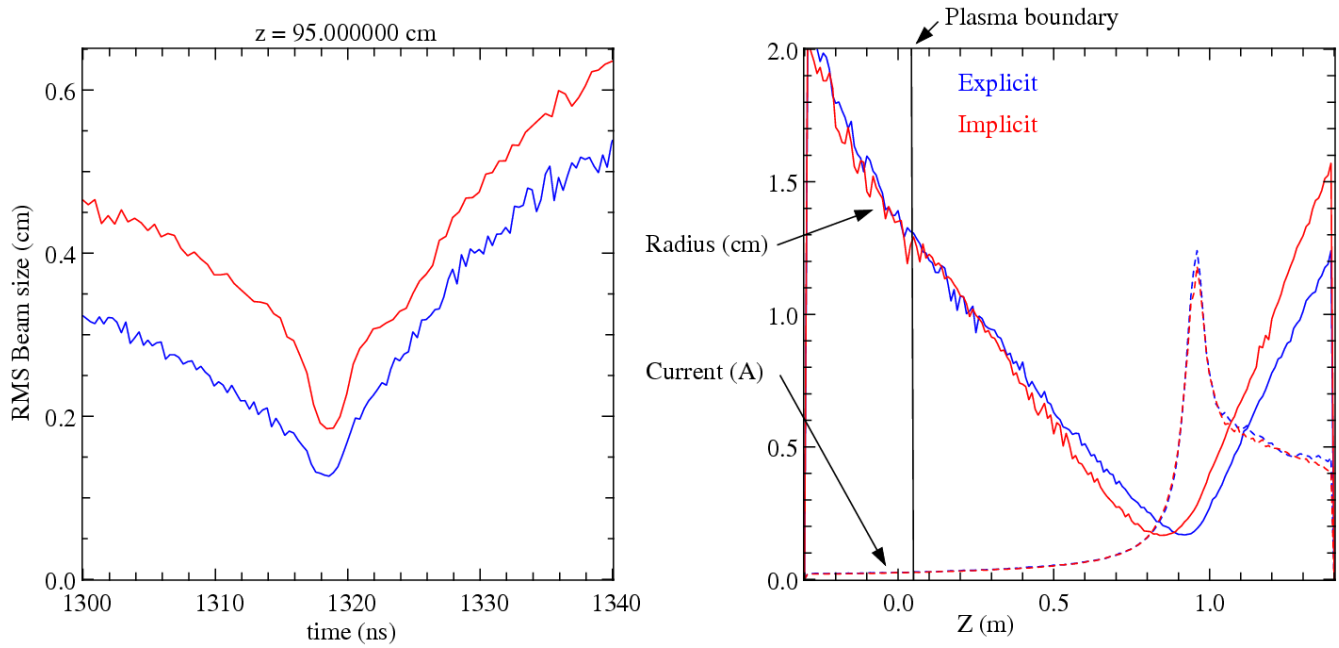

(a) (b)

Figure 7. Comparison of implicit (red) and explicit (blue) Warp runs: (a) beam radius at the z-location of peak axial compression; (b) beam radius and current vs. position. The current here is the maximum observed as the beam passes through each position. The radius at each z is the RMS radius at the time of maximum current at that z. As can be seen, the implicit case focuses ~5 cm before the explicit case.

E. Drift-Lorentz mover tests

We have performed two types of tests of the implicit version of the drift-Lorentz mover. The first consists of following a set of single-electron orbits in a quadrupole magnetic field, with a superposed fixed potential well representing the effect of an ion beam. This tests the validity of the particle advance scheme independent of the field solution. We have previously shown that the explicit drift-Lorentz mover with $\omega_{ce}\Delta t > 1$ reproduces very well the results of full-ion-dynamics simulations with small (cyclotron-period-resolving) timesteps[iii]. Here, we show comparison of orbits from the drift-Lorentz mover using the standard (explicit) advance scheme and the implicit advance scheme. In both cases a large timestep, $\omega_{ce}\Delta t = 16$, is used. Referring to Fig. 8, traces a and e in each plot correspond to orbits launched on field lines which pass very close to the magnetic-field null; traces b and d are for electrons launched somewhat further from the null field lines; and trace c is for an electron that does not get near the null. The frequency of the stair-steps is a measure of the bounce frequency of the electrons, and the average slope is a measure of the drift rate (electric plus magnetic). Hence, it can be seen that orbits a and e are non-adiabatic (magnetic moment is not conserved, so the bounce frequency changes), b and d are marginally adiabatic, and c is adiabatic. From the figure, it can be seen that the explicit and implicit movers agree very well for adiabatic and marginally adiabatic orbits, and disagree as expected for nonadiabatic orbits (which are chaotic).

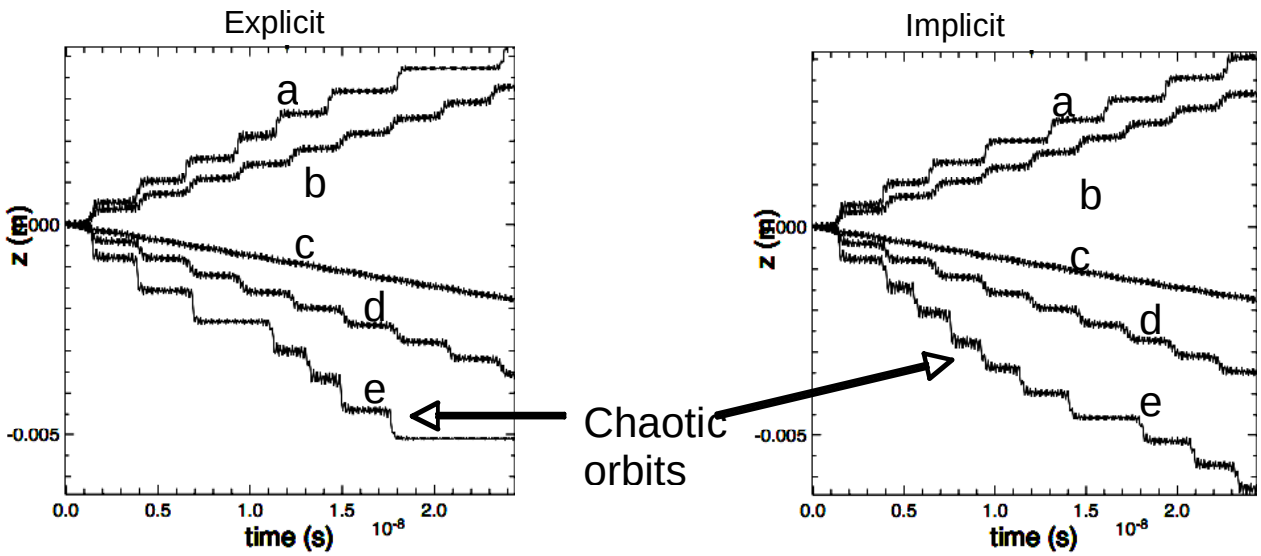

Figure 8. Comparison of single-particle orbits obtained with explicit and implicit movers for electrons in quadrupole field with superposed electrostatic potential (see text).

The second class of tests derives from of the Ion-Temperature Gradient (ITG) instability case discussed in Sec. IIB. These results are still very preliminary. The simulation setup is the same as in Sec. IIIB, except that the parameters are chosen to be those of the classic ITG test case of Lee and Tang[vii], and as in that work, an artificial drift is added to simulate the effect of temperature and density gradients. The parameters differing from Sec. IIIB are: $T_i = T_e = 4.8$ KeV, $B_y = 1$ T, $B_z = 0.01\ B_y = 0.01$ T, density $n = 9.7\times10^{18}$ m^{-3}, and $\Delta t = 2.8\times10^{-8}$ s. For these parameters, $\omega_{pe}\Delta t = \omega_{ce}\Delta t = 272$, while $\omega_{pi}\Delta t = 27.2$ and $\omega_{ci}\Delta t = 2.72$. Also, the ion thermal gyroradius equals the cell size. We have performed the following simulations: (a) uniform plasma, (b) density gradient only, with $\kappa_n \equiv (\rho/n)\ dn/dx = 0.05$; and (c) density and temperature gradients, with $\kappa_n = 0.05$ and $\kappa_T \equiv (\rho/T)\ dT/dx = 0.05$. According to theory, cases (a) and (b) should be stable, while case (c) should be strongly unstable. This theoretical result is borne out by our (preliminary) simulations, as shown in the Fig. 9: for cases (a) and (b) the fluctuating potential decreases from seed level and remains small; for case (c) the r.m.s. fluctuating potential rises almost immediately to a level two orders of magnitude higher than for cases (a) and (b), roughly consistent with the saturation levels observed by Lee and Tang. Our initial seed is apparently too noisy to observe the linear growth phase. (In contrast to the studies reported in our earlier work with an adiabatic (Boltzmann) electron-response model[iv], where linear growth was observed, the electrons in this case are kinetic, and the simulation, as expected, is noisier).

We have also rerun the uniform-plasma case with the pure Boris direct-implicit scheme (no interpolation with drift kinetics). In this case, the Courant condition is violated for both ions and electrons, and the simulation is numerically unstable. It is for this reason that modified parameters were adopted for the tests of Sec. IIIB.

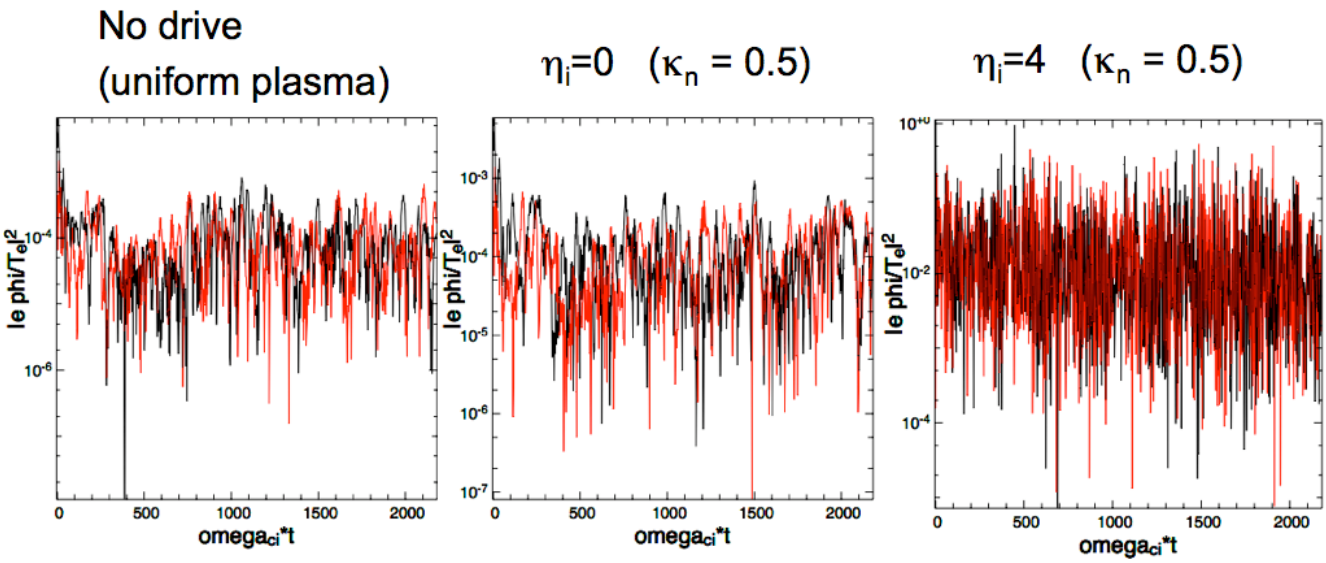

Figure 9. Plots of square of the fluctuation amplitude in first two Fourier modes for the ITG test problem, for the case of a uniform plasma (left), a density gradient but uniform temperature (center), and a physically unstable case with both density and temperature gradients (right). Black and red curves correspond to x,z wavenumbers (1,1) and (1,-1), respectively.

F. Wave excitation by a beam propagating through plasma

Using parameters similar to those of earlier simulations that were carried out with the EDPIC and LSP codes, we performed explicit electromagnetic Warp simulations in 2-D (x,z) geometry of wave excitations by a beam propagating through plasma. See Fig. 1 of Ref. viii; a relevant subset, which used EDPIC, is reproduced below. In these simulations, a hard-edged, elliptical, "frozen" (rigid) beam propagates at constant velocity $v_z = 0.5c$ through an initially cold neutral plasma of initial density n_0. The beam has a flat-top density profile of $n_b = n_0/2$, and an elliptical shape of length $l = 15\ c/\omega_p$ and diameter $d = l/10$, where ω_p is the electron plasma frequency. A snapshot of the electron density taken from the Warp calculation is shown in Fig. 10. In agreement with the simulation results and analysis of the reference, waves with a wavenumber approximately $2\omega_p/v_z$ are generated in the plasma by the beam's electrostatic field. These are of high amplitude inside the beam, due to the sharp beam edges. We also note visible differences with the earlier results, which we attribute to the use of different detailed parameters (the EDPIC run used $l = 20\ c/\omega_p$, and boundary conditions also differed).

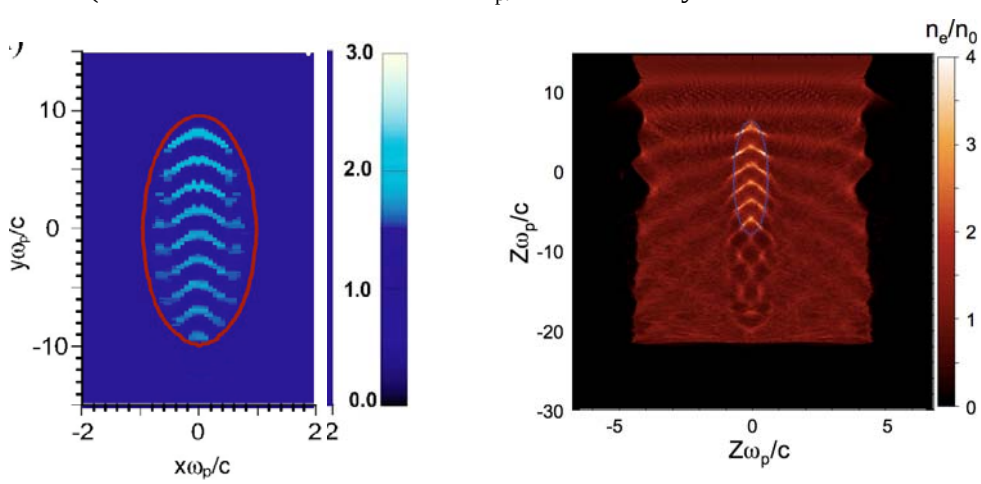

Figure 10. Electron density n_e from 2-D "slab" (x,z) explicit electromagnetic PIC simulations of a rigid beam (thin outline) propagating through a neutral plasma, using EDPIC (left) and Warp (right). Parameters differ.

REFERENCES

i. A. Friedman, A. B. Langdon, and B. I. Cohen, "A Direct Method for Implicit Particle-in-Cell Simulation," *Comments on Plasma Physics and Controlled Fusion* **6**, 225 (1981).

ii. C. K. Birdsall and A. B. Langdon, *Plasma Physics via Computer Simulation*, Section 14.13.

iii. R. H. Cohen, A. Friedman, M. Kireeff Covo, *et al.*, *Phys. Plasmas* **12**, 056708 (2005).

iv. R. H. Cohen, A. Friedman, D. P. Grote and J.-L. Vay, *Nucl. Inst. Meth.* **A 577**, 52 (2007).

v. D. R. Welch, *et al*., paper THPAN086, *Proc. 2007 Particle Accelerator Conference (PAC07)*, Albuquerque, New Mexico, USA (2007).

vi. A. B. Sefkow, Ph.D. Thesis, Princeton University, 2007, available at: http://www.pppl.gov/~asefkow/thesis .

vii. W. W. Lee and W. M. Tang, *Phys. Fluids* **31**, 612 (1986).

viii. I. D. Kaganovitch, *et al.*, "Nonlinear plasma waves excitation by intense ion beams in background plasma," *Phys. Plasmas* **11**, 3546 (2004).