LA-UR-15-25379
Approved for public release; distribution is unlimited.

Title:	Fully implicit Particle-in-cell algorithms for multiscale plasma simulation
Author(s):	Chacon, Luis
Intended for:	Plasma Energization: Exchanges Between Fluid and Kinetic Scales, 2015-05-04/2015-05-06 (Los Alamos, New Mexico, United States) Web
Issued:	2015-07-16

Fully implicit Particle-in-cell algorithms for multiscale plasma simulation

L. Chacón

Applied Mathematics and Plasma Physics Group
Theoretical Division
Los Alamos National Laboratory
P.O. Box 1663, Los Alamos, NM 87544

Collaborators:
G. Chen, W. Taitano, D. A. Knoll (LANL)
D. C. Barnes (Coronado Consulting)
and the (now dispersed) CoCoMANS team

Plasma Energization: Exchanges between Fluid and Kinetic Scales
May 4-6, 2015, LANL

Work funded by the LANL LDRD program

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

Outline

- Particle-in-cell (PIC) methods for plasmas
- Explicit vs. implicit PIC
- 1D ES implicit PIC
  - Charge and energy conservation
  - Moment-based acceleration
- Generalization to Multi-D EM PIC: Vlasov-Darwin model
  - Review and motivation for Darwin model
  - Conservation properties (energy, charge, and canonical momenta)
  - Numerical benchmarks

Introduction

Luis Chacon, chacon@lanl.gov

Kinetic Plasma Simulation

- A fully ionized collisionless plasma: ions, electrons, and electromagnetic fields
- **Challenge**: integrate electron-ion-field kinetic system on an ion time-scale and a system length scale while retaining electron kinetic effects accurately.

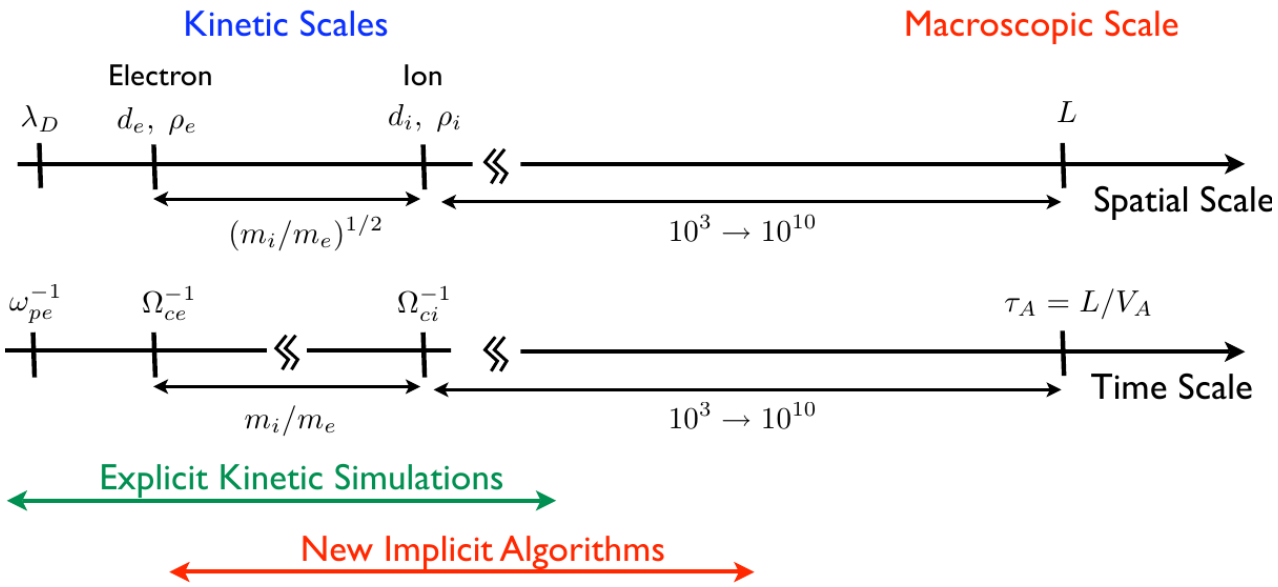

(We are developing a new implicit algorithm for long-term, system-scale simulations.)

- Problem features a **hierarchical description**:
  - How to design a multiscale algorithm?
  - How to respect conservation laws, and constraints?

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacon, chacon@lanl.gov

Particle-in-cell (PIC) methods for kinetic plasma simulation

$$\partial_t f + \mathbf{v}\cdot\nabla f + \frac{\mathbf{F}}{m}\cdot\nabla_v f = 0$$

- Lagrangian solution by the **method of characteristics**:

$$f(\mathbf{x},\mathbf{v},t) = f_0\left(\mathbf{x} - \int_0^t dt\mathbf{v}, \mathbf{v} - \frac{1}{m}\int_0^t dt\mathbf{F}\right)\ ;\ \mathbf{x}(t=0)=\mathbf{x}_0\ ;\ \mathbf{v}(t=0)=\mathbf{v}_0$$

- PIC approach follows characteristics employing **macroparticles** (volumes in phase space)

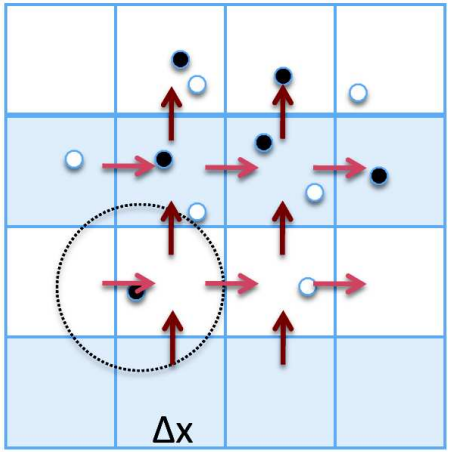

$$f(\mathbf{x},\mathbf{v},t) = \textstyle\sum_p \delta(\mathbf{x}-\mathbf{x}_p)\delta(\mathbf{v}-\mathbf{v}_p)$$

$$\begin{aligned} \dot{\mathbf{x}}_p &= \mathbf{v}_p \\ \dot{\mathbf{v}}_p &= \frac{q_p}{m_p}(\mathbf{E}+\mathbf{v}\times\mathbf{B}) \end{aligned}$$

$$\begin{aligned} \partial_t\mathbf{B} + \nabla\times\mathbf{E} &= 0 \\ -\mu_0\epsilon_0\partial_t\mathbf{E} + \nabla\times\mathbf{B} &= \mu_0\mathbf{j} \\ \nabla\cdot\mathbf{B} &= 0 \\ \nabla\cdot\mathbf{E} &= \frac{\rho}{\epsilon_0} \end{aligned}$$

$$\delta(\mathbf{x}-\mathbf{x}_p) \longrightarrow S(\mathbf{x}-\mathbf{x}_p)\ ;\ E_p = \sum_i E_i S(x_i - x_p)\ ;\ j_i = \sum_p j_p S(x_i - x_p)$$

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

State-of-the-art *classical* PIC algorithm is explicit

- Classical explicit PIC: "leap-frogs" particle positions and velocities, field-solve at position update:

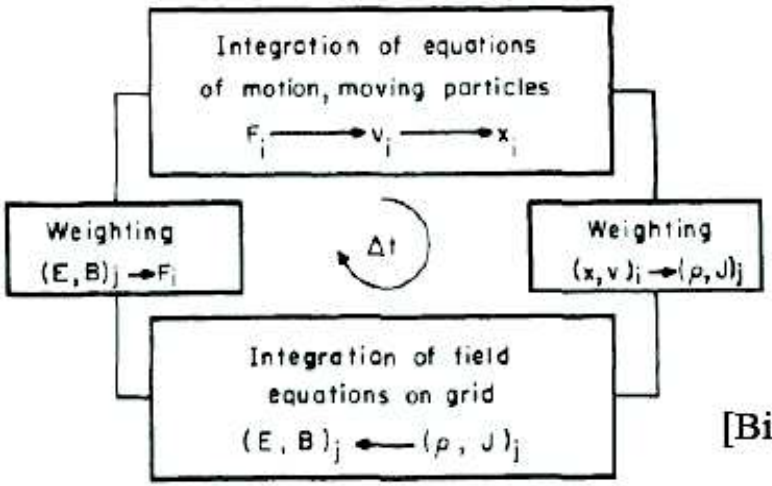

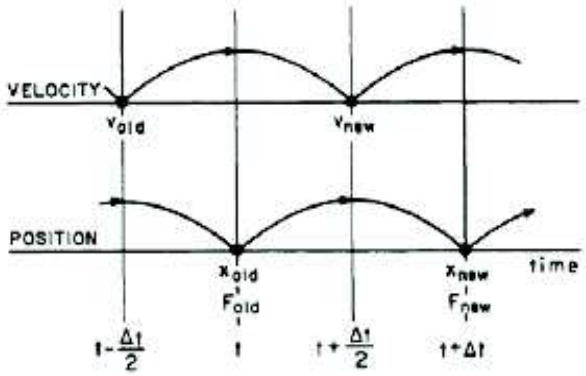

[Birdsall and Langdon, Plasma physics via computer simulation]

- Implementation is straightforward, but...
- Performance limitations:
  - CFL-type instability: $\min(\omega_{pe}\Delta t < 1, c\Delta t < \Delta x)$. **Minimum temporal resolution**
  - Finite-grid instability: $\Delta x < \lambda_{Debye}$. **Minimum spatial resolution**
  - Memory bound: challenging for efficient use of modern computer architectures.
- Accuracy limitations:
  - **Lack of energy conservation**, problematic for long-time-scale simulations
- To remove the stability/accuracy constraints of explicit methods, we consider implicit methods.

Implicit PIC methods

- Exploration of implicit PIC started in the 1980s
  - Implicit moment method [1]
  - Direct implicit method [2]
- Early approaches used linearized, semi-implicit formulations:
  - Lack of nonlinear convergence
  - Particle orbit accuracy (particle and fields integrated in lock-step)
  - Inconsistencies between particles and moments
  - Inaccuracies! $\rightarrow$Plasma self-heating/cooling [3]
- Our approach: nonlinear implicit PIC
  - Enforcing nonlinear convergence; complete consistency between particles, moments, and fields.
  - Allowing stable and robust integrations with large Δt and Δx ($2nd\ order$ accurate).
  - Ensuring exact global energy conservation and local charge conservation properties.
  - Allowing adaptivity in both time and space without loss of the conservation properties.
  - Allowing particle subcycling $\rightarrow$ high operational intensities ($compute\ bound$).
  - Allowing **fluid preconditioning** to $accelerate$ the $iterative$ kinetic solver!

[1] Mason, R. J. (1981), Brackbill, J. U., and Forslund, D. W. (1982)
[2] Friedman, A., Langdon, A. B. and Cohen, B. I.(1981)
[3] Cohen, B. I., Langdon, A. B., Hewett, D. W., and Procassini, R. J. (1989)

Fully implicit PIC: 1D electrostatic PIC

Chen et al, JCP 2011, 2012, 2013; Taitano et al, SISC (2013)

Luis Chacon, chacon@lanl.gov

Fully implicit PIC formulation (at first glance)

- A fully implicit formulation couples particles and fields non-trivially (integro-differential PDE):

$$\begin{aligned} \frac{f^{n+1}-f^n}{\Delta t} + \mathbf{v}\cdot\nabla\frac{f^{n+1}+f^n}{2} \quad - \quad & \frac{q}{m}\nabla\frac{\Phi^{n+1}+\Phi^n}{2}\cdot\nabla_{\mathbf{v}}\frac{f^{n+1}+f^n}{2} = 0 \\ \nabla^2\Phi^{n+1} \quad = \quad & \int d\mathbf{v}\, f^{n+1}(\mathbf{x},\mathbf{v},t) \end{aligned}$$

- In PIC, f^{n+1} is sampled by a large collection of particles in phase space, $\{\mathbf{x},\mathbf{v}\}_p^{n+1}$.
  - There are N_p particles, each particle requiring $2\times d$ equations ($d\rightarrow$dimensions),
  - Field requires N_g equations, one per grid point.
- If implemented naively, an impractically large algebraic system of equations results:

$$\boxed{\mathbf{G}(\{\mathbf{x},\mathbf{v}\}_p^{n+1},\{\Phi^{n+1}\}_g)=0} \;\rightarrow\; \dim(\mathbf{G}) = 2dN_p + N_g$$

  - No current computing mainframe can afford the memory requirements
  - Algorithmic issues are showstoppers (e.g., how to precondition it?)
- An alternative strategy exists: nonlinear elimination (**particle enslavement**)

Particle enslavement (nonlinear elimination)

- Full residual $\mathbf{G}(\{x,v\}_p,\{\Phi\}_g)=0$ is impractical to implement
- Alternative: nonlinearly eliminate particle quantities **so that they are not dependent variables**:
  - Formally, particle equations of motion are functionals of the electrostatic potential:

$$x_p^{n+1}=x_p[\Phi^{n+1}]\ ;\ v_p^{n+1}=v_p[\Phi^{n+1}]$$

$$\mathbf{G}(\mathbf{x}_p{}^{n+1},\mathbf{v}_p{}^{n+1},\mathbf{\Phi}^{n+1})=\mathbf{G}(\mathbf{x}[\mathbf{\Phi}^{n+1}],\mathbf{v}[\mathbf{\Phi}^{n+1}],\mathbf{\Phi}^{n+1})=\tilde{\mathbf{G}}(\mathbf{\Phi}^{n+1})$$

Nonlinear residual can be *unambiguously* formulated in terms of electrostatic potential only!

- JFNK storage requirements are dramatically decreased, making it tractable:
  - Nonlinear solver storage requirements $\propto N_g$, **comparable to a fluid simulation**
  - Particle quantities $\Rightarrow$ auxiliary variables: only a **single copy of particle population** needs to be maintained in memory throughout the nonlinear iteration

Energy-conserving (EC) Vlasov-Ampère discretization

- Fully implicit Crank-Nicolson time discretization:

$$\varepsilon_0 \frac{E_i^{n+1} - E_i^n}{\Delta t} + \sum_p q_p v_p^{n+1/2} S(x_i - x_p^{n+1/2}) = 0$$

$$\frac{x_p^{n+1} - x_p^n}{\Delta t} = \frac{v_p^{n+1} + v_p^n}{2}$$

$$\frac{v_p^{n+1} - v_p^n}{\Delta t} = \frac{q_p}{m_p} \sum_i \frac{E_i^n + E_i^{n+1}}{2} S(x_i - x_p^{n+1/2})$$

In time:
- centered, 2nd order;
- implicit;
- unconditionally stable;
- **non-dissipative.**

- C-N enforces energy conservation to numerical round-off:

$$\sum_p \frac{m_p}{2}(v_p^{n+1} + v_p^n)(v_p^{n+1} - v_p^n) = -\sum_i \varepsilon_0 \frac{E_i^{n+1} - E_i^n}{\Delta t} \frac{E_i^{n+1} + E_i^n}{2} \Rightarrow \sum_p \frac{1}{2} m_p v_p^2 + \sum_i \frac{1}{2} \varepsilon_0 E_i^2 = \text{const}$$

- As a result, the formulation does not suffer from finite-grid instabilities (normal mode analysis)
  - ⇨ Unconstrained spatial resolution: $\Delta x \not< \lambda_D$!!
- Energy conservation is only realized when particles and fields are nonlinearly converged:
  - ⇨ Requires a tight nonlinear tolerance

Jacobian-Free Newton-Krylov Methods

- After spatial and temporal discretization $\Rightarrow$ a large set of nonlinear equations: $\boxed{\vec{G}(\vec{x}^{n+1}) = \vec{0}}$
- Converging nonlinear couplings requires iteration: Newton-Raphson method:

$$\left.\frac{\partial \vec{G}}{\partial \vec{x}}\right|_k \delta\vec{x}_k = -\vec{G}(\vec{x}_k)$$

- Jacobian linear systems result, which require a linear solver $\Rightarrow$ Krylov subspace methods (GMRES)
  - Only require matrix-vector products to proceed.
  - Jacobian-vector product can be computed Jacobian-free (**CRITICAL**: no need to form Jacobian matrix):

$$\left(\frac{\partial \vec{G}}{\partial \vec{x}}\right)_k \vec{y} = J_k\vec{y} = \lim_{\epsilon\to 0}\frac{\vec{G}(\vec{x}_k + \epsilon\vec{y}) - \vec{G}(\vec{x}_k)}{\epsilon}$$

  - Krylov methods can be easily preconditioned: $P_k^{-1} \sim J_k^{-1}$

$$J_k P_k^{-1} P_k \delta\vec{x} = -\vec{G}_k$$

We will explore suitable preconditioning strategies later in this talk.

Algorithmic implementation details

- The nonlinear residual formulation $\mathbf{G}(E^{n+1})$ based on Vlasov-Ampere formulation is as follows:
  1. Input E (given by JFNK iterative method)
  2. Move particles (i.e., find $x_p[E]$, $v_p[E]$ by solving equations of motion)
     (a) Requires inner (local) nonlinear iteration: Picard (not stiff)
     (b) Can be as complicated as we desire (substepping, adaptivity, etc)
  3. Compute moments (current)
  4. Form Vlasov-Ampere equation residual
  5. return
- Because particle move is performed within function evaluation, we have much freedom.
- We can explore improvements in particle mover to ensure **long-term accuracy!**
  - **Particle substepping and orbit averaging** (ensures orbit accuracy and preserves exact energy conservation)
  - **Exact charge conservation strategy** (a new charge-conserving particle mover)
  - **Orbit adaptivity** (to improve momentum conservation)

Particle orbit substepping

- In applications of interest, field time-scale (Δt) and orbit time-scale ($\Delta\tau$) can be well separated
  - Fields evolve *slowly* (dynamical time scale, Δt)
  - Particle orbits may still undergo *rapid change* ($\Delta\tau \ll \Delta t$)
- Particle orbits need to be resolved to **avoid large orbit integration errors**

Accurate orbit integration requires particle substepping!

- Field does not change appreciably: time-averaged value over long time scale is sufficient

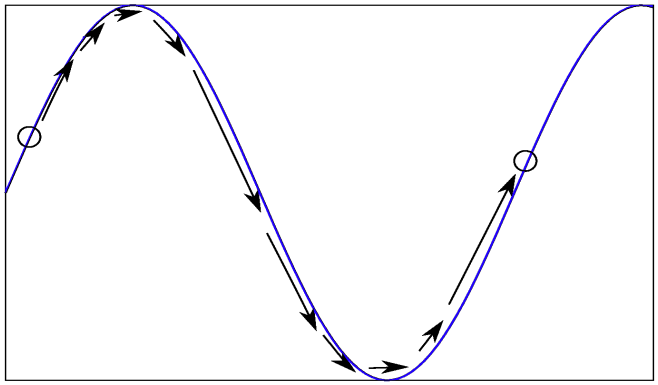

$$\frac{x_p^{\nu+1} - x_p^{\nu}}{\Delta\tau} = v_p^{\nu+1/2}$$

$$\frac{v_p^{\nu+1} - v_p^{\nu}}{\Delta\tau} = \sum_i \underbrace{\frac{E_i^{n+1} + E_i^{n}}{2}}_{\text{slow}} S(x_i - x_p^{\nu+1/2})$$

Energy conservation and orbit averaging

- **Particle substepping breaks energy conservation**.
- Energy conservation theorem can be recovered by orbit averaging Ampère's law:

$$\epsilon_0 \partial_t E + j = \langle j \rangle \quad , \quad \frac{1}{\Delta t}\int_t^{t+\Delta t} d\tau [\cdots] \Rightarrow \epsilon_0 \frac{E^{n+1} - E^n}{\Delta t} + \bar{j} = \langle \bar{j} \rangle$$

- Orbit-averaged current is found as:

$$\bar{j} = \frac{1}{\Delta t}\int_t^{t+\Delta t} d\tau\, j \approx \frac{1}{\Delta t}\sum_p \sum_{\nu=1}^{N_\nu} q_p v_p S(x - x_p)\Delta\tau^\nu$$

- With these definitions, exact energy conservation is recovered:

$$\sum_p \sum_\nu \frac{m_p}{2}(v_p^{\nu+1} + v_p^\nu)(v_p^{\nu+1} - v_p^\nu) = -\sum_i \epsilon_0 \frac{E^{n+1} - E^n}{\Delta t}\frac{E_i^{n+1} + E_i^n}{2}$$

$$\Rightarrow \sum_p \frac{1}{2} m_p v_p^2 + \sum_i \frac{1}{2}\epsilon_0 E_i^2 = \text{const.}$$

Luis Chacon, chacon@lanl.gov

Exact charge conservation: charge-conserving particle mover

- Local charge conservation $\partial_t \rho + \nabla \cdot \mathbf{j} = 0$ is generally violated in the discrete.
- Local charge conservation is essential to ensure long-term accuracy (can be derived independently from both Vlasov and Maxwell equations; "glues" them together).
- B-spline interpolation ensure charge conservation within cell boundaries; charge conservation broken when particles cross cell boundaries.
  - **Standard strategy** based on current redistribution when particle crosses boundary. [Buneman 1968, Morse and Nielson, 1971]
  - Current redistribution breaks energy conservation. **Need a new strategy**.

Here, charge conservation is enforced by stopping particles at cell boundaries.

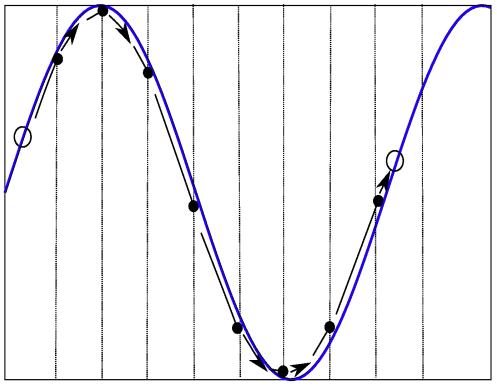

$$\left.\begin{aligned} \rho_{i+\frac{1}{2}} &= \sum_p q_p \frac{S_m(x - x_{i+\frac{1}{2}})}{\Delta x} \\ j_i &= \sum_p q_p v_p \frac{S_{m-1}(x - x_i)}{\Delta x} \\ S'_m(x) &= \frac{S_{m-1}(x+\frac{\Delta x}{2}) - S_{m-1}(x-\frac{\Delta x}{2})}{\Delta x} \end{aligned}\right\} \overset{(m=1,2)}{\Longrightarrow} \left[\partial_t \rho + \nabla \cdot \mathbf{j} = 0\right]_{i+\frac{1}{2}}^{n+\frac{1}{2}} = 0$$

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacon, chacon@lanl.gov

Improved momentum conservation: particle orbit adaptivity

- EC/CC PIC algorithm does not enforce momentum conservation exactly.
  - ⇨ **Controlling error** in momentum conservation is **crucial** for long-term accuracy
- Orbit integration errors can significantly affect momentum conservation: **particle tunneling**

- Adaptive orbit integration can be effective in suppressing particle tunneling and thus improve momentum conservation
- **Approach**: find $\Delta\tau$ to control local truncation error. Second order estimator gives:

$$\Delta\tau \leq \sqrt{12\epsilon_r \frac{m_p}{q_p} \left|\frac{dE}{dx}\right|_p^{-1}}$$

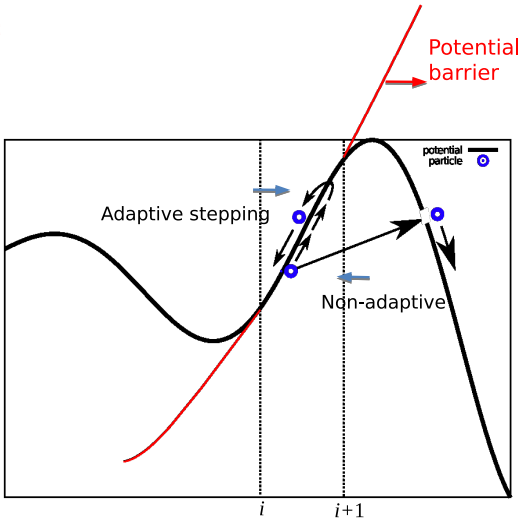

- Electric field gradient is estimated from cell-based gradient: $\left.\frac{\partial E}{\partial x}\right|_p \approx \frac{E_{i+1}-E_i}{\Delta x}$. **Provides potential barrier!**

- **Particle is stopped at cell boundaries** to ensure charge conservation.

Luis Chacon, chacon@lanl.gov

Ion acoustic shock wave

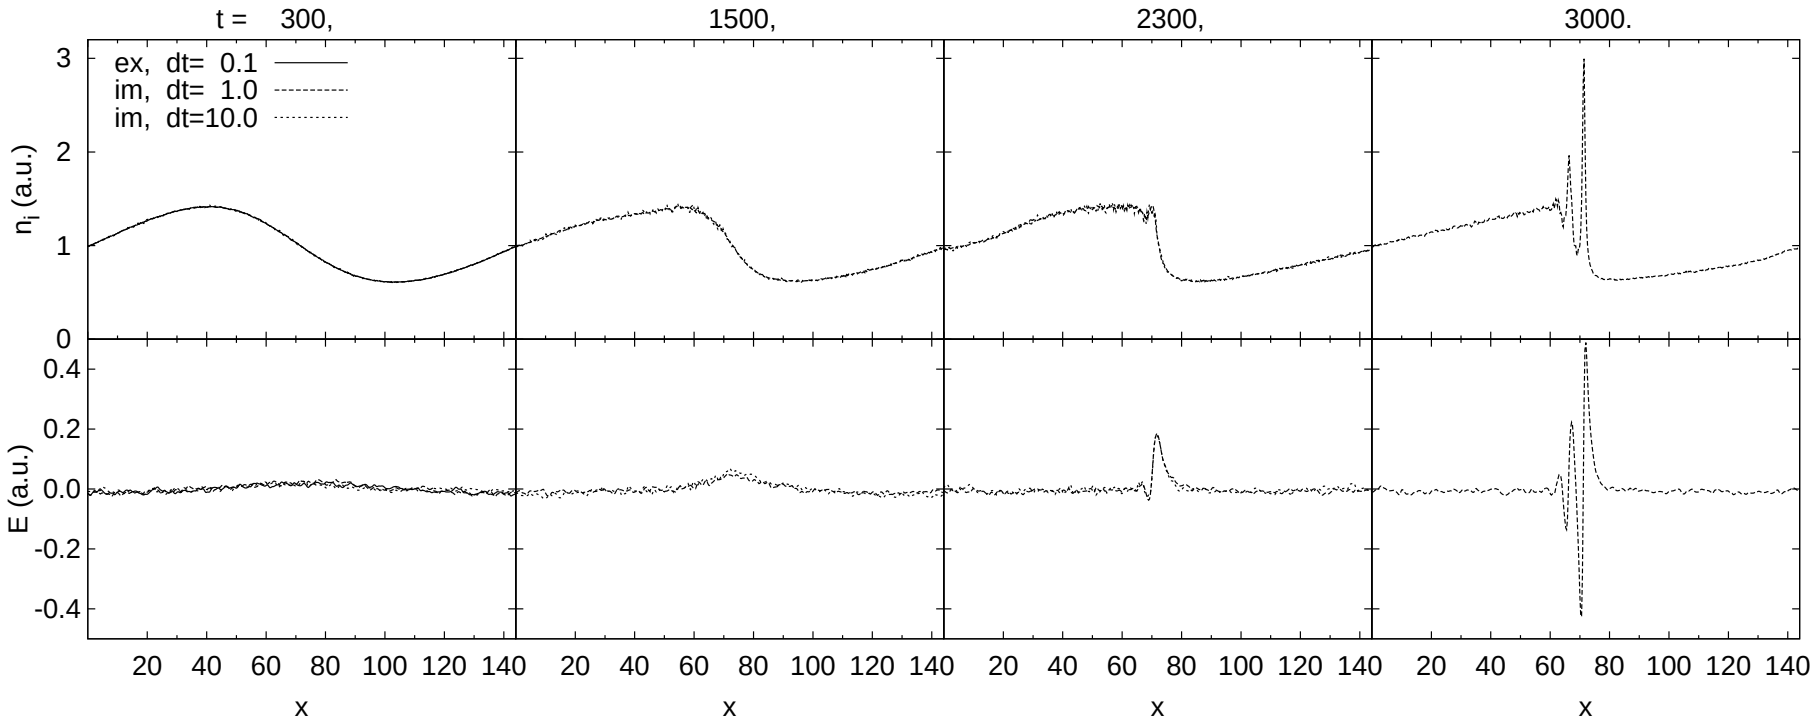

- Propagating IAW with perturbation level $\epsilon = 0.4$, with 4000 particles/cell.
- Realistic mass ratio ($m_i/m_e = 2000$).
- Shock wave length scale$\sim$Debye length.

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

Ion acoustic shock wave test

non-uniform grid spacing

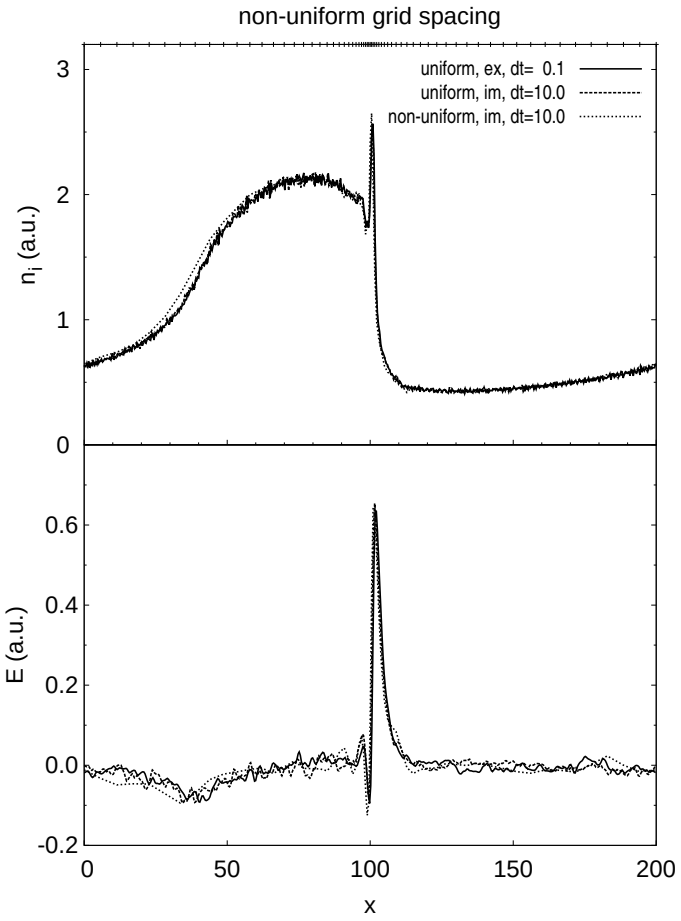

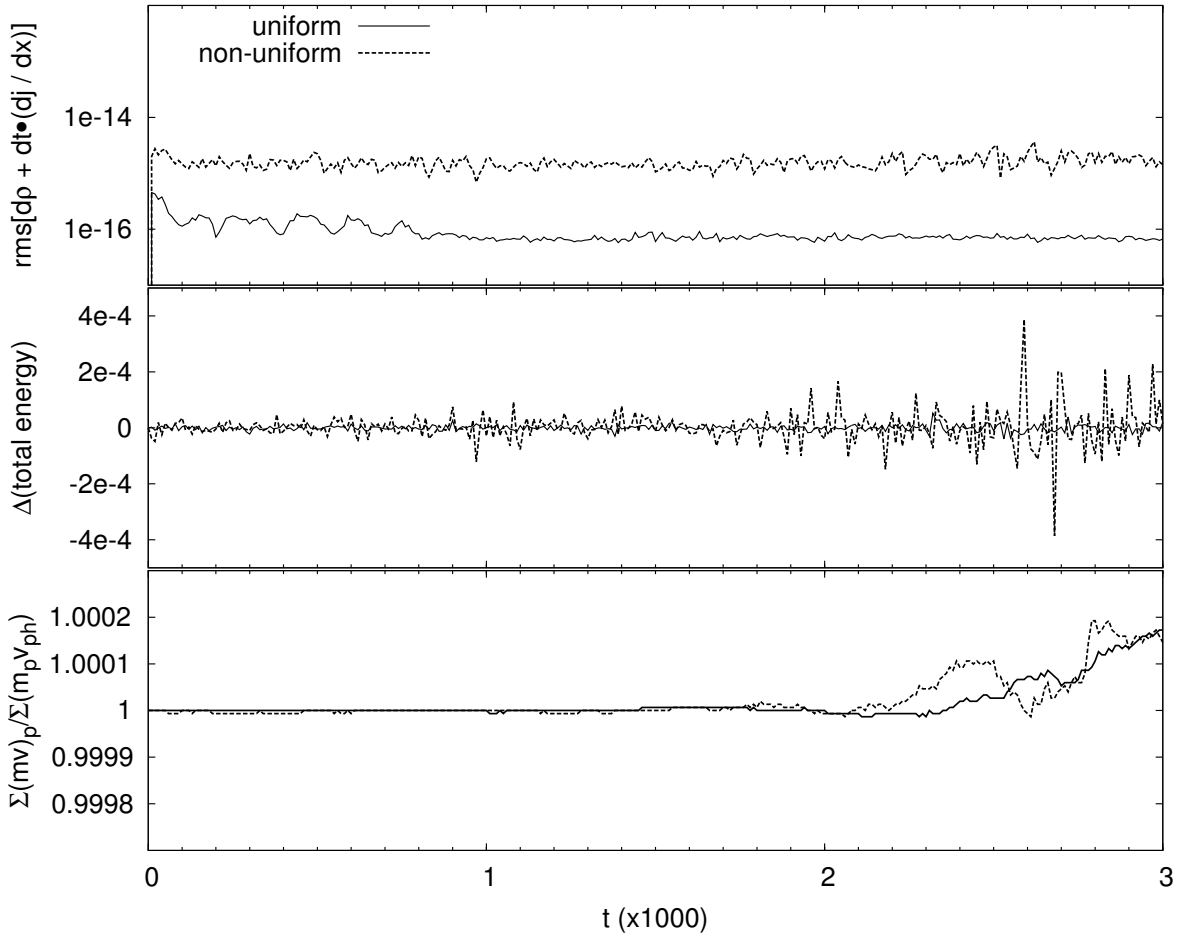

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

Comparison against Implicit Moment Method

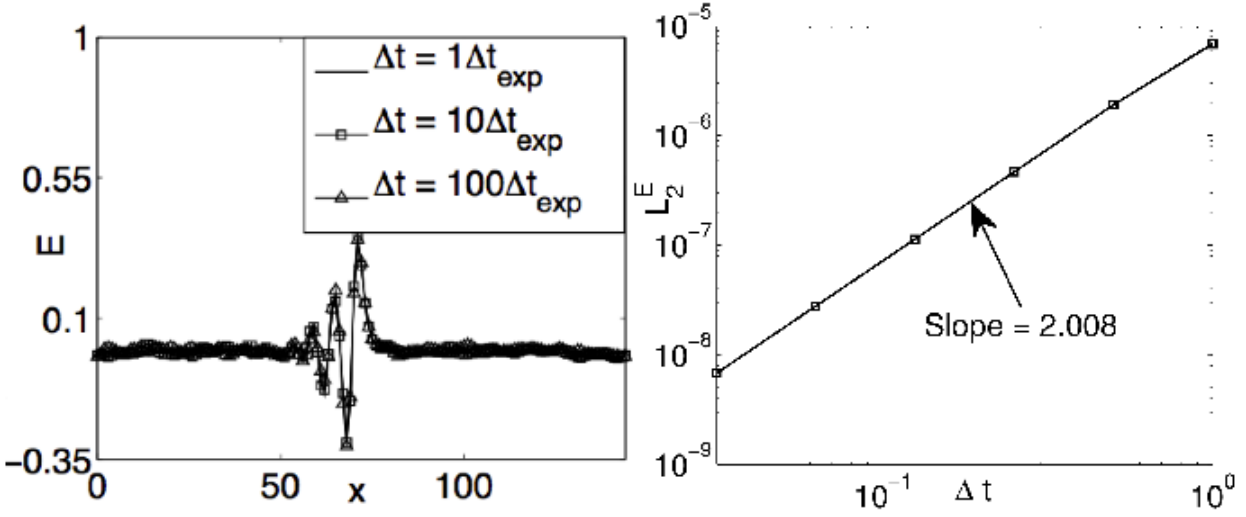

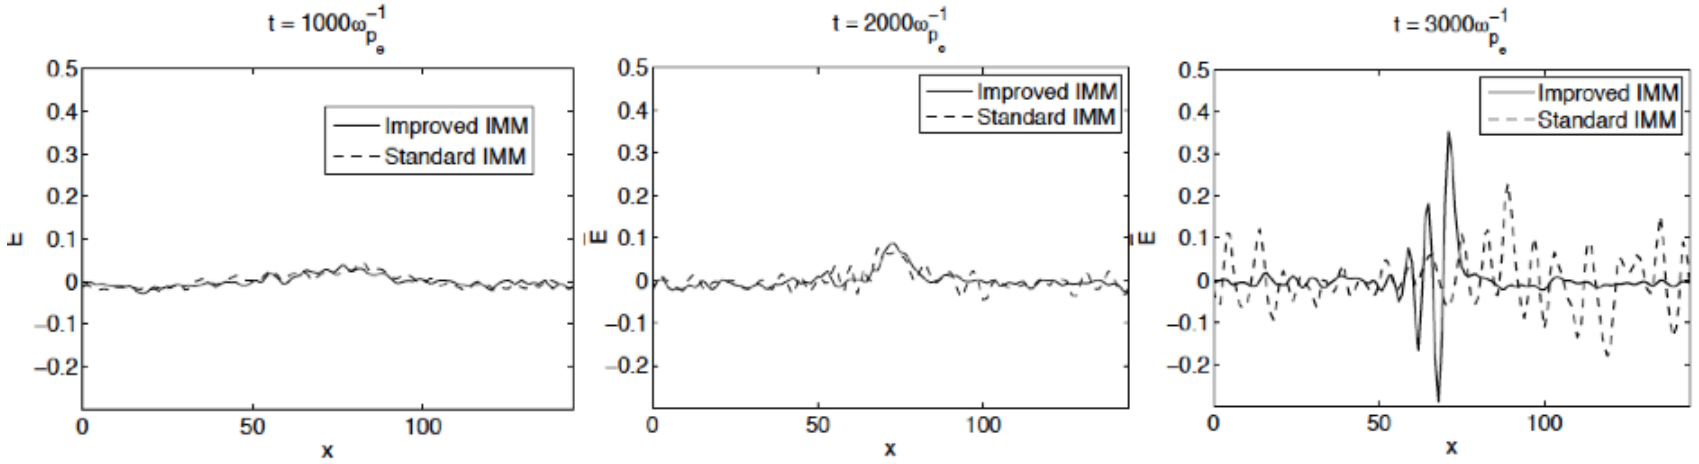

Luis Chacon, chacon@lanl.gov

Moment-based acceleration of fully implicit kinetic solver

Chen et al., JCP (2014)

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

CPU gain potential of implicit PIC vs. explicit PIC

- Back-of-the-envelope estimate of CPU gain:

$$CPU \sim \left(\frac{T}{\Delta t}\right)\left(\frac{L}{\Delta x}\right)^d n_p C^{solver} \quad ; \quad \frac{C^{imp}}{C^{ex}} \sim N_{FE}\frac{\Delta t_{imp}}{\Delta \tau_{imp}} \quad ; \quad \frac{CPU_{ex}}{CPU_{imp}} \sim \left(\frac{\Delta x_{imp}}{\Delta x_{ex}}\right)^d \frac{\Delta \tau_{imp}}{\Delta t_{ex}}\frac{1}{N_{FE}}$$

- Using reasonable estimates:

$$\begin{aligned}
\Delta \tau_{imp} &\sim \min\left[0.1\frac{\Delta x_{imp}}{v_{th}}, \Delta t_{imp}\right] \\
\Delta t_{imp} &\sim 0.1\omega_{pi}^{-1} \\
\Delta t_{exp} &\sim 0.1\omega_{pe}^{-1} \\
k\Delta x_{imp} &\sim 0.2 \\
\Delta x_{ex} &\sim \lambda_D
\end{aligned}$$

$$\frac{CPU_{ex}}{CPU_{imp}} \sim \frac{1}{(k\lambda_D)^d}\frac{1}{N_{FE}}\min\left[\frac{1}{k\lambda_D}, \sqrt{\frac{m_i}{m_e}}\right]$$

- CPU speedup is:
  - Better for realistic mass ratios and increased dimensionality!
  - Limited by solver performance N_{FE} (preconditioning!)

Moment-based acceleration of fully kinetic simulations

- Particle elimination $\Rightarrow$ nonlinear residual is formulated in terms of fields/moments ONLY: $\mathbf{G}(E)$
- Within JFNK, preconditioner ONLY needs to provide field/moment update:

$$\delta E \approx -P^{-1}\mathbf{G}$$

Premise of acceleration: obtain δE from a fluid model using current particle distribution for closure.

- We begin with corresponding fluid nonlinear model:

$$\begin{aligned}
\partial_t n_\alpha &= -\nabla \cdot \mathbf{\Gamma}_\alpha \\
m_\alpha \left[\partial_t \mathbf{\Gamma}_\alpha + \nabla \cdot \left(\frac{1}{n_\alpha} \mathbf{\Gamma}_\alpha \mathbf{\Gamma}_\alpha \right) \right] &= q_\alpha n_\alpha \mathbf{E} + \nabla \cdot \left(n_\alpha \left(\frac{\mathbf{\Pi}_\alpha}{n_\alpha} \right)_p \right) \\
\epsilon_0 \partial_t \mathbf{E} &= \sum_\alpha q_\alpha \mathbf{\Gamma}_\alpha
\end{aligned}$$

Luis Chacon, chacon@lanl.gov

Moment-based acceleration of fully kinetic simulations (cont.)

➤ We formulate *approximate* linearized fluid equations (neglect linear temperature response):

$$\frac{\delta n_\alpha}{\Delta t} = -\nabla \cdot \delta \mathbf{\Gamma}_\alpha$$

$$m_\alpha \frac{\delta \mathbf{\Gamma}_\alpha}{\Delta t} \approx q_\alpha (\delta n_\alpha \mathbf{E} + n_{\alpha,p} \delta \mathbf{E}) + \nabla \cdot \left(\left(\frac{\mathbf{\Pi}_\alpha}{n_\alpha} \right)_p \delta n_\alpha \right)$$

$$\epsilon_0 \delta \mathbf{E} = \Delta t \left[\sum_\alpha q_\alpha \delta \mathbf{\Gamma}_\alpha - \mathbf{G}(\mathbf{E}) \right]$$

δE can be obtained from Newton state $\mathbf{E}$, Newton residual $\mathbf{G}(\mathbf{E})$, and particle closures $\mathbf{\Pi}_{\alpha,p}$ and $n_{\alpha,p}$

Luis Chacon, chacon@lanl.gov

Preconditioner performance: CPU scaling

$$\frac{CPU_{ex}}{CPU_{imp}} \sim \frac{1}{(k\lambda_D)^d}\frac{1}{N_{FE}}\min\left[\frac{1}{k\lambda_D}, \sqrt{\frac{m_i}{m_e}}\right]$$

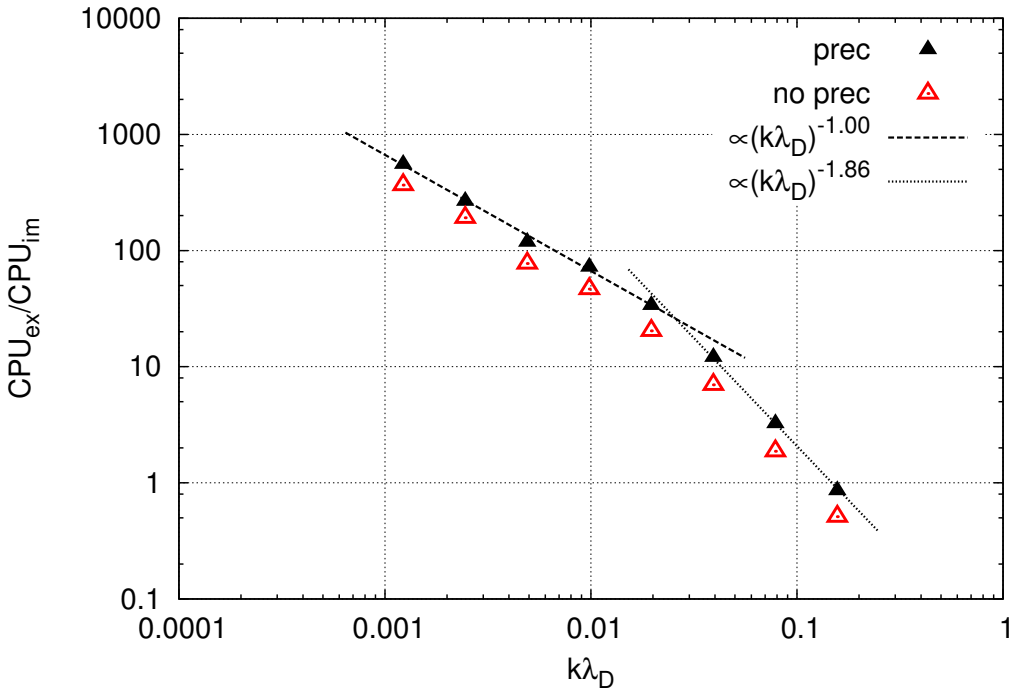

Transition occurs at $k\lambda_D \sim \sqrt{\frac{m_e}{m_i}} \sim 0.025$, as predicted

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacon, chacon@lanl.gov

Electromagnetic PIC: non-radiative Darwin formulation

Chen et al, CPC 2014, 2015 (submitted)

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

Implicit time-stepping and numerical dispersion of light wave[4]

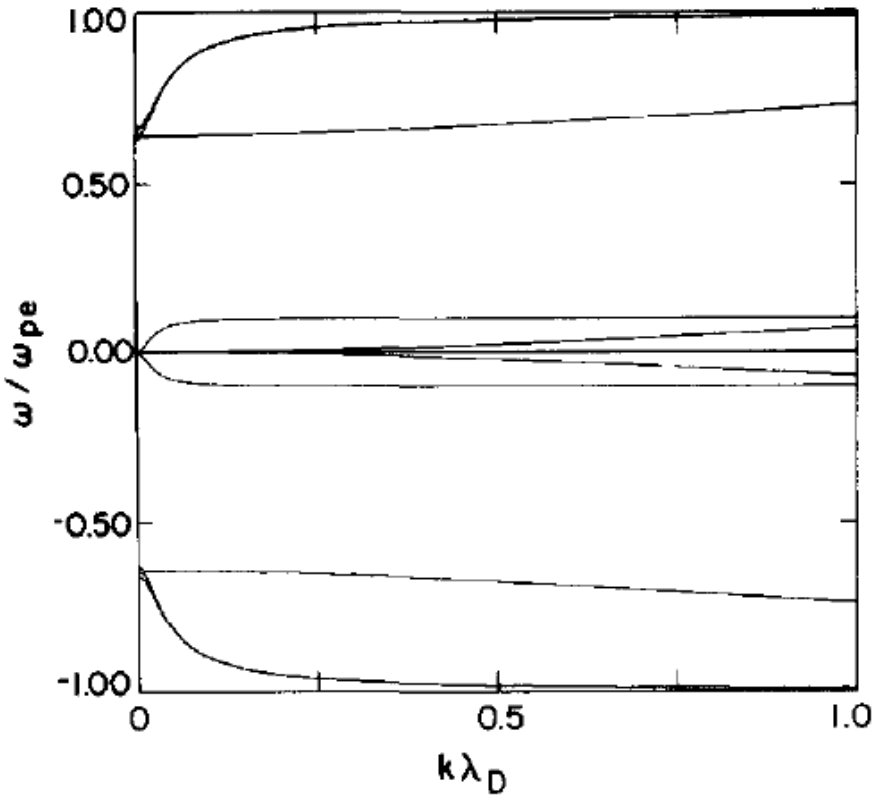

- **Implicit time-stepping** introduces numerical dispersion to light wave
  - ⇨ Artificially decreases speed of light
- Fast particles become supra-luminal numerically ⇒ light-wave excitation and radiative noise

[4] Brackbill, J. U., and Forslund, D. W. (1982)

Luis Chacon, chacon@lanl.gov

Light wave excitation and radiative noise in real simulations

- If one keeps light wave with exact energy conservation, one gets **enhanced numerical noise due to numerical Cherenkov radiation (or possibly instability).**
- Noise-control requires numerical dissipation

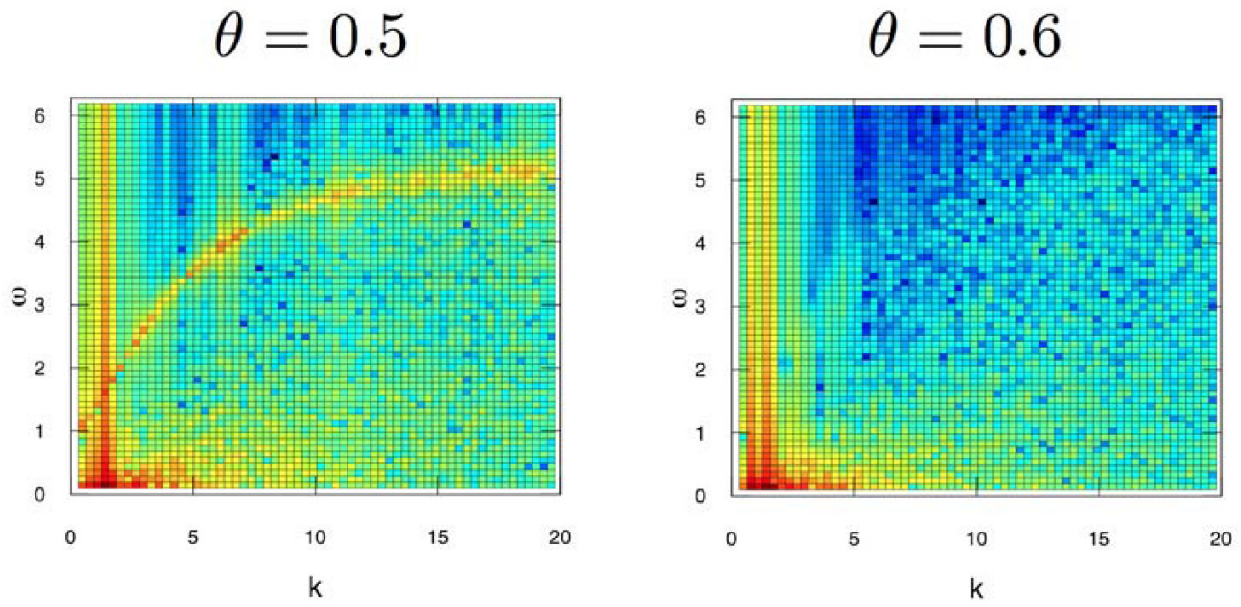

Figure 1: Fourier spectrum for Weibel instability. [Markidis and Lapenta, JCP 2011].

- Numerical dissipation breaks energy conservation
- **Solution**: analytically remove light-wave when relativistic effects are not important

Darwin model formulation (potential form)

- Darwin model is formal $\mathcal{O}(v/c)^2$ approximation to Maxwell's equations[5]
  - Analytically removes light-wave *while preserving charge separation effects*
- Begin with Maxwell's equations:

$$\begin{aligned} \partial_t \mathbf{B} + \nabla \times \mathbf{E} &= \mathbf{0}, & (1)\\ \frac{1}{c^2}\partial_t \mathbf{E} + \mu_0 \mathbf{j} &= \nabla \times \mathbf{B}, & (2)\\ \nabla \cdot \mathbf{E} &= \rho/\epsilon_0, & (3)\\ \nabla \cdot \mathbf{B} &= 0. & (4) \end{aligned}$$

- Consider potentials ϕ, $\mathbf{A}$ such that:

$$\begin{aligned} (4) \Rightarrow \mathbf{B} &= \nabla \times \mathbf{A}, \\ (1) \Rightarrow \mathbf{E} &= -\nabla\phi - \partial_t \mathbf{A}. \end{aligned}$$

- In the Coulomb gauge ($\nabla \cdot \mathbf{A} = 0$), taking $c \to \infty$ in transverse displacement current:

$$(2) \Rightarrow \cancel{\frac{1}{c^2}\frac{\partial^2 A}{\partial t^2}} - \nabla^2 \mathbf{A} = \mu_0 \left[\mathbf{j} - \epsilon_0 \nabla \partial_t \phi\right],$$

$$\nabla \cdot (2) \Rightarrow \epsilon_0 \nabla^2 \partial_t \phi = \nabla \cdot \mathbf{j}.$$

[5] Krause and Morrison (2007)

Luis Chacon, chacon@lanl.gov

Darwin model formulation (cont.)

- **Full Darwin system**:

$$-\frac{1}{\mu_0}\nabla^2\mathbf{A} = \mathbf{j} - \epsilon_0\partial_t\nabla\phi, \tag{5}$$

$$\epsilon_0\partial_t\nabla^2\phi = \nabla\cdot\mathbf{j}. \tag{6}$$

$$\nabla\cdot\mathbf{A} = 0 \tag{7}$$

$$\nabla^2\phi = -\rho/\epsilon_0 \tag{8}$$

- Enforcing involutions (Eqs. 7, 8) is critical for accuracy.
- Careful discretization allows one to **imply** involutions, rather than solving for them:
  - $\nabla\cdot\mathbf{A} = 0$ implied by Eqs. 5, 6 and careful boundary conditions: $\boxed{\nabla^2(\nabla\cdot\mathbf{A}) = 0}$
  - $\nabla^2\phi = -\rho/\epsilon_0$ implied by Eq. 6 and exact PIC charge conservation: $\boxed{\partial_t\rho + \nabla\cdot\mathbf{j} = 0}$
- **Minimal Darwin formulation** ($\mathbf{j}$ obtained from particles):

$$\nabla^2\chi = \nabla\cdot\mathbf{j},$$

$$-\nabla^2\mathbf{A} = \mu_0\left[\mathbf{j} - \nabla\chi\right],$$

$$\epsilon_0\partial_t\phi = \chi.$$

Numerical integration of Vlasov-Darwin in multi-D

- Work directly with potential formulation, avoiding explicit involutions
  - Spatial discretization on a Yee mesh
  - Automatic enforcement of Coulomb gauge ($\nabla \cdot \mathbf{A} = 0$) and Gauss' law ($\nabla^2 \phi = -\rho/\epsilon_0$)
  - **NO divergence cleaning needed!**
- Fully implicit, fully nonlinear time stepping (Crank-Nicolson)
  - Particles are nonlinearly enslaved, subcycled, time-adapted (implicit Boris push)
  - Exact local charge conservation (implies Gauss' law)
  - Exact global energy conservation
  - Exact conservation of canonical momenta in ignorable directions

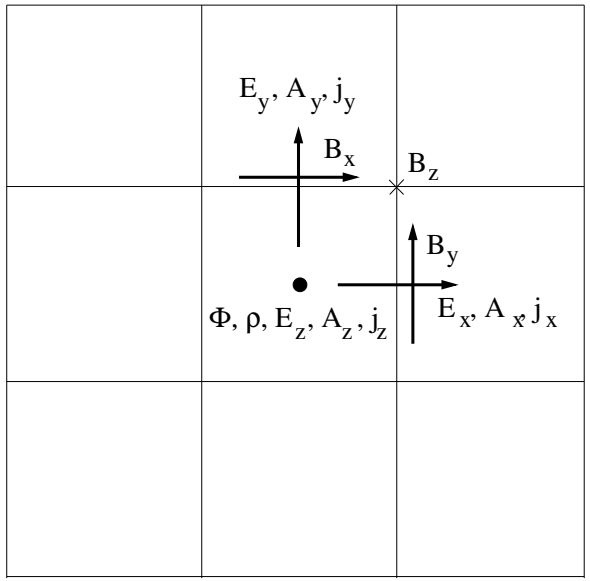

Luis Chacon, chacon@lanl.gov

CPU speedup potential of EM implicit PIC vs. explicit PIC

- Back-of-the-envelope estimate of CPU gain:

$$\frac{CPU_{ex}}{CPU_{imp}} \sim \left(\frac{\Delta x_{imp}}{\Delta x_{ex}}\right)^d \frac{\Delta \tau_{imp}}{\Delta t_{ex}} \frac{1}{N_{FE}}$$

$$\begin{aligned}
\Delta \tau_{imp} &\sim \min\left[0.1\frac{\Delta x_{imp}}{v_{th,e}}, 0.1\omega_{ce}^{-1}, \Delta t_{imp}\right] \\
\Delta t_{imp} &\sim 0.1\omega_{pi}^{-1} \\
\Delta t_{exp} &\sim \frac{\Delta x_{exp}}{c} \\
k\Delta x_{imp} &\sim 0.2 \\
\Delta x_{ex} &\sim \lambda_D
\end{aligned}$$

$$\frac{CPU_{ex}}{CPU_{imp}} \sim \frac{1}{(k\lambda_D)^d} \frac{c}{v_{th,e}} \frac{1}{N_{FE}} \min\left[\frac{1}{k\lambda_D}, \frac{c}{v_A}\sqrt{\frac{m_e}{m_i}}, \sqrt{\frac{m_i}{m_e}}\right]$$

- CPU speedup is:
  - Impacted by electron-ion mass ratio, how close electrons are to relativistic speeds.
- Again, **key is to minimize** N_{FE}.
  - We have developed a very effective moment-based preconditioner.

EM preconditioner summary

- We have developed a fluid preconditioner that takes into account both ion and electron linear responses:
  - ⇨ Proper asymptotic behavior:
    - Large domain sizes ($L \gg d_i$): recover Hall MHD and MHD current responses
    - Small electron-to-ion mass ratios, $m_e/m_i \ll 1$
  - ⇨ Effective for $\omega_{pe} > \omega_{ce}$, i.e., weakly to moderately magnetized plasmas
    - $\frac{m_e}{m_i} > \left(\frac{v_A}{c}\right)^2$, i.e., it limits achievable mass ratio for fixed magnetic field
    - Could be overcome with proper model for gyroviscous linear response
    - If strongly magnetized regimes are of interest, then $\Delta t_{imp} \lesssim 0.1\omega_{ce}^{-1}$, and:

$$\frac{CPU_{ex}}{CPU_{imp}} \sim \frac{1}{(k\lambda_D)^d} \frac{c}{v_{th,e}} \frac{1}{N_{FE}}$$

Still strong potential for algorithmic acceleration.

Luis Chacon, chacon@lanl.gov

1D Electron Weibel instability

- Isotropic ions, bi-Maxwellian electrons

$$m_i/m_e = 1836,\ T_{e\perp}/T_{e\parallel} = 16,\ N_{e,i}\text{=128,000},\ L = 2\pi c/\omega_{pe},\ N_g\text{=32}.$$

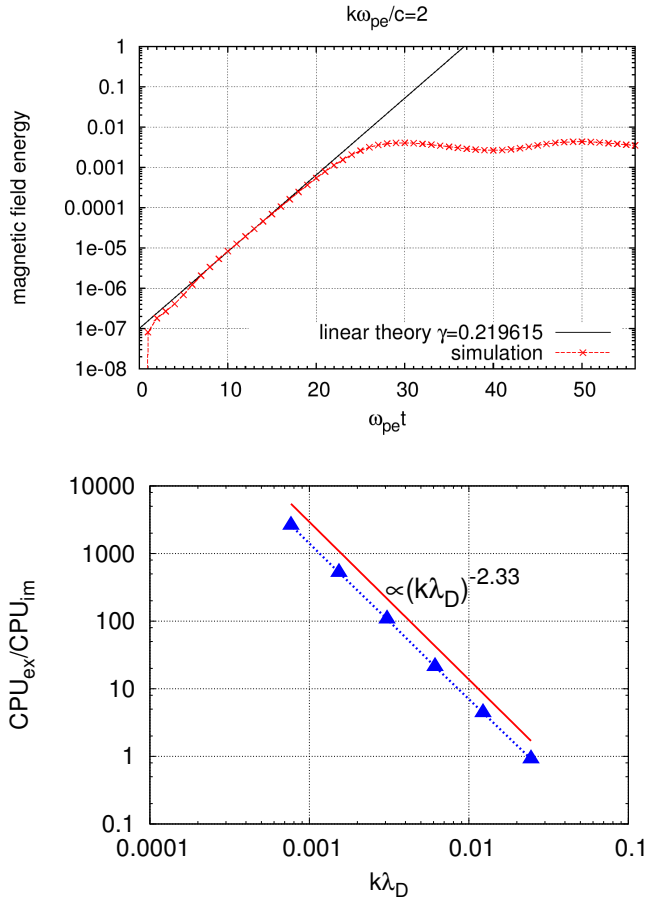

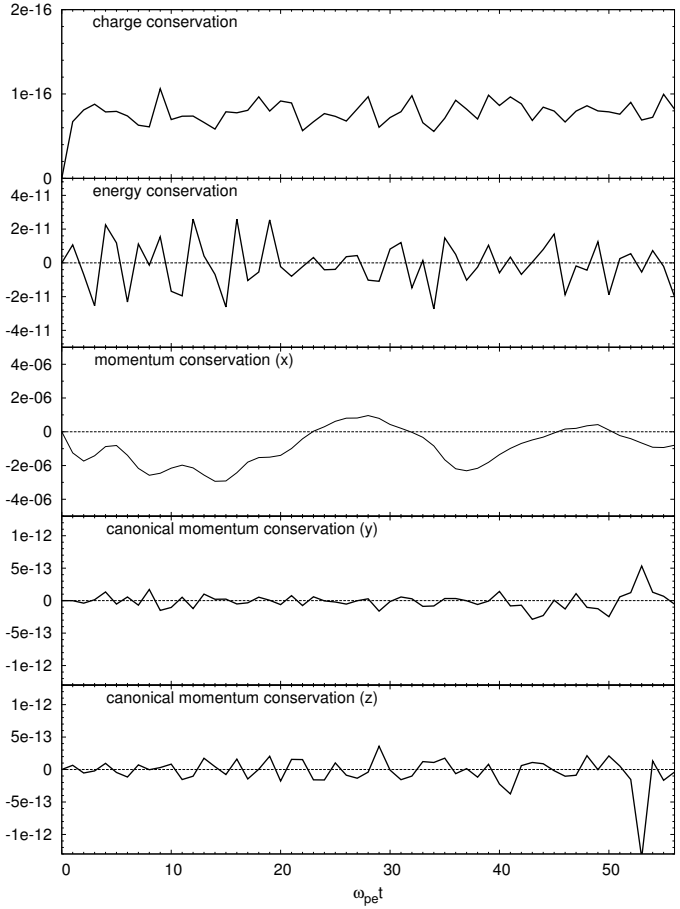

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacon, chacon@lanl.gov

1D Electron Weibel instability (cont.)

- Numerical demonstration of second-order accuracy in time

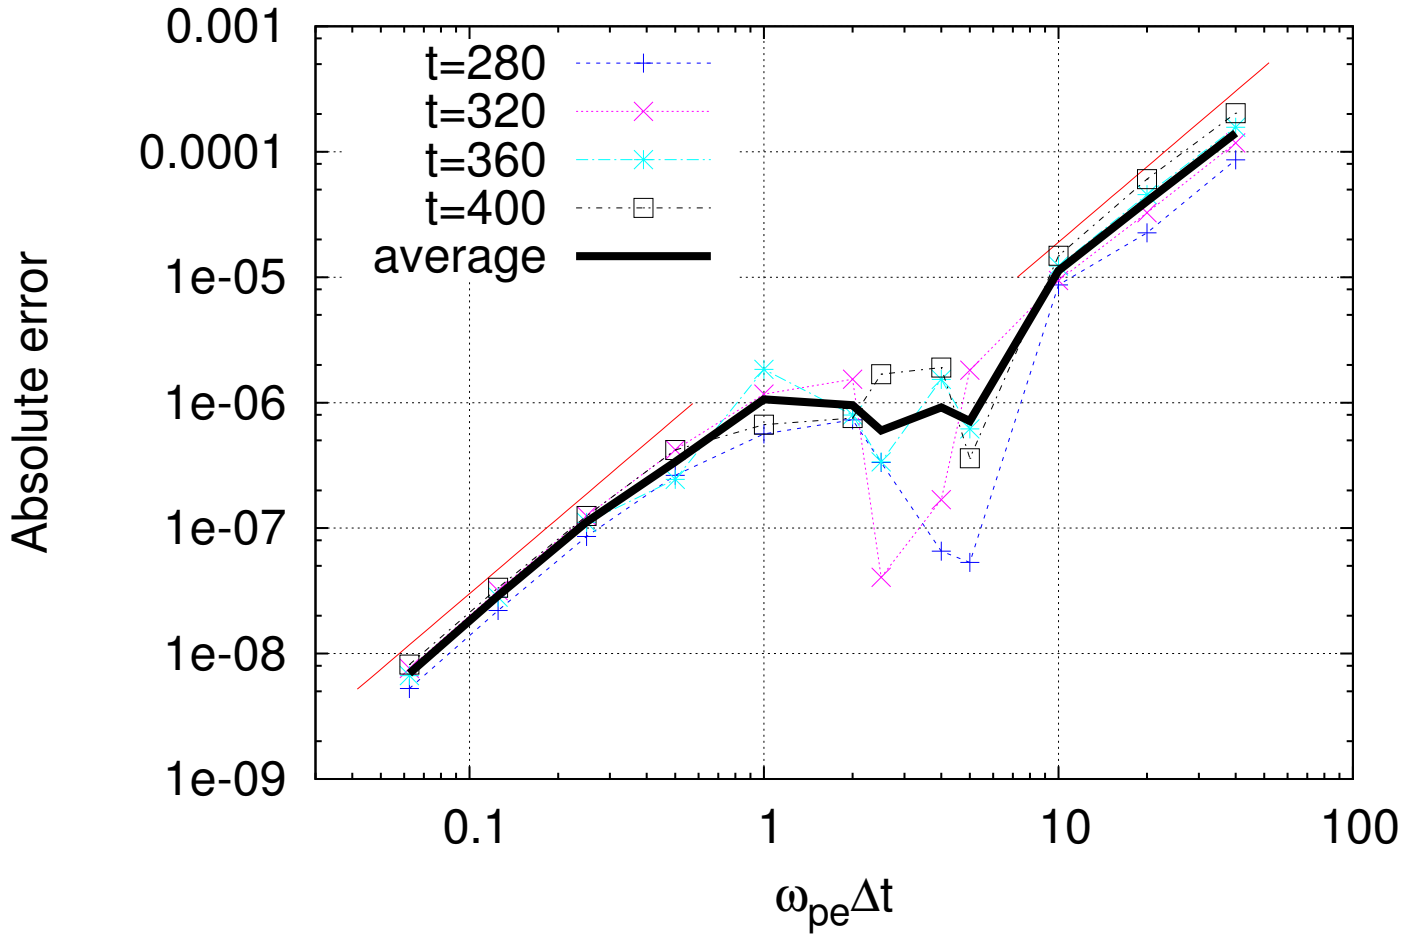

Luis Chacon, chacon@lanl.gov

1D Kinetic Alfven Wave[6]

$$B_0 = 0.00778\ ;\ m_i/m_e = 1836\ ;\ k\lambda_{De} = 0.003\ ;\ v_{th,e}/c = 0.07\ ;\ (\beta_e = 0.1)$$

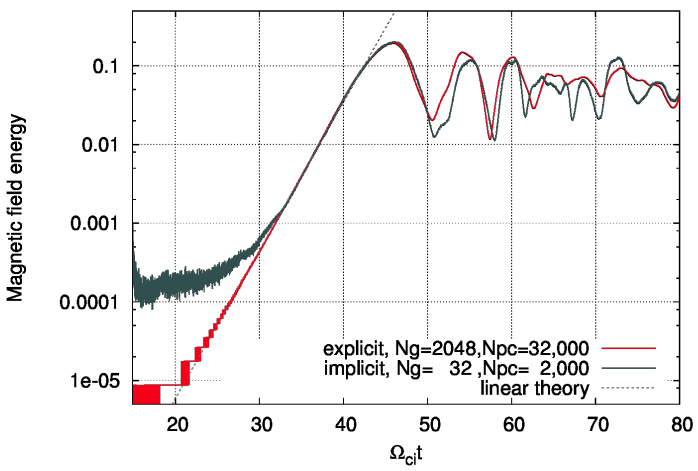

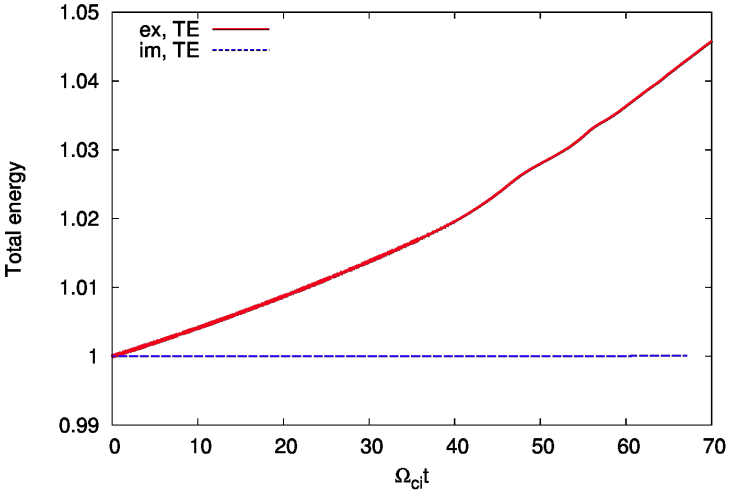

- Explicit PIC (Daughton's NPIC):
  - 2048 mesh points, 32,000 pcles/cell (overkill for this problem), **5% energy error**
  - 500 CPUs x 24 hr, 7×10^6 time steps
- Implicit PIC:
  - 32 mesh points, 2,000 pcles/cell ($1000\times$ fewer particles), 10^{-4}**% energy error**
  - 16 CPUs x 29 hr, 1.3×10^5 time steps, $N_{FE} \sim 30$ $(r_{tol} = 10^{-6})$
- CPU speedup ~ 26 ($\times 100$ in 2D, $\times 10^4$ in 3D)

[6] Yin et al., POP **14** (2007)

2D Weibel instability

$m_i/m_e = 1836$, $T_{e\perp}/T_{e\parallel} = 9$, N_{pc}=2000, $L = \pi\, d_e \times \pi\, d_e$, $N_g = 32 \times 32$

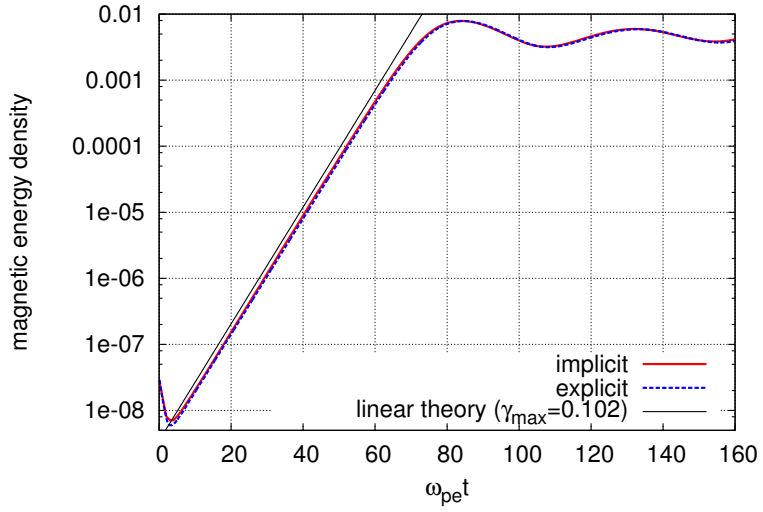

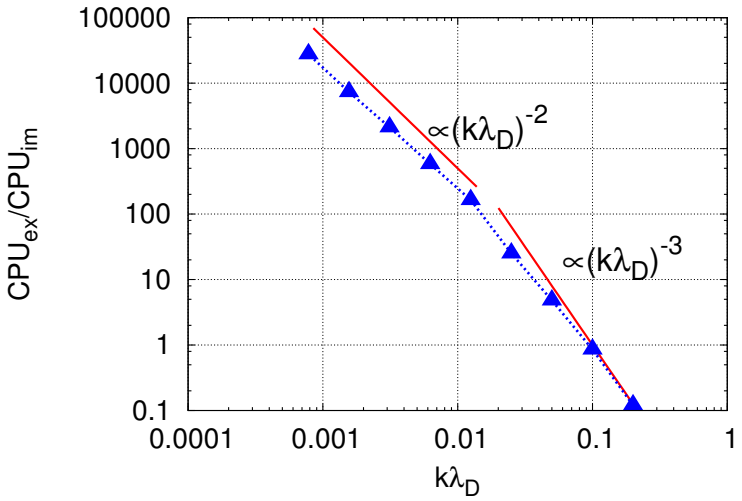

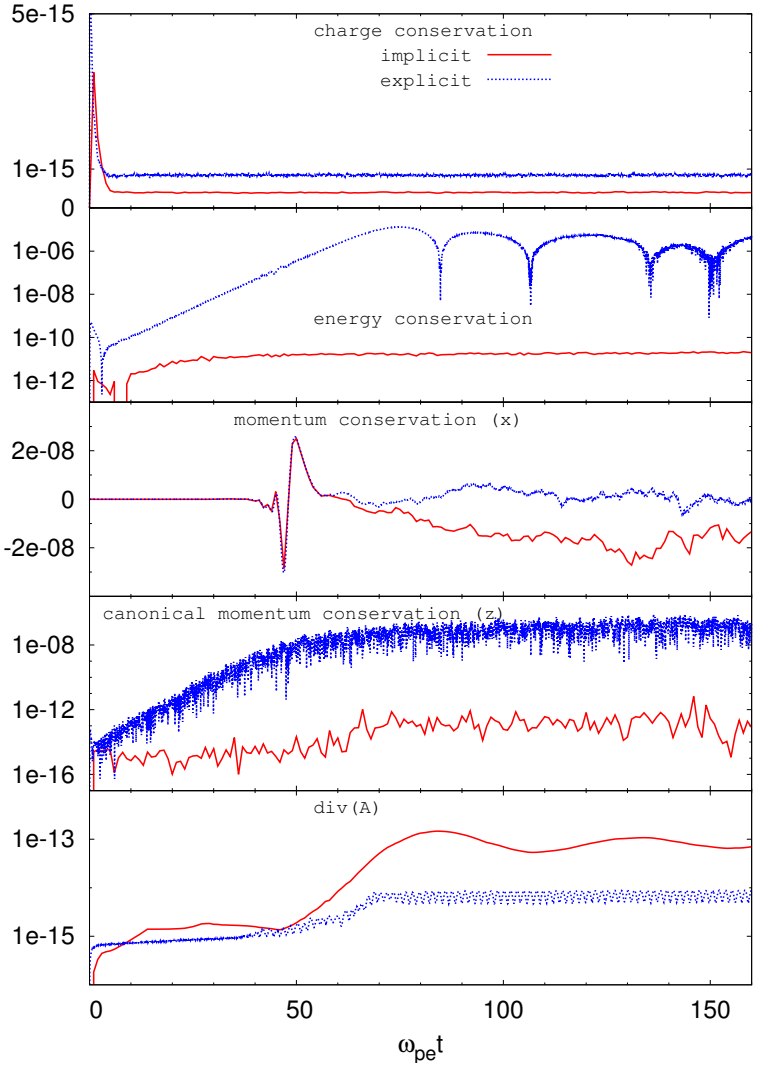

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

2D Weibel instability (non-uniform mesh)

$$L_x \times L_y = 22\,d_e \times 22\,d_e,\ N_{pc} = 800,\ \Delta t = 0.1\omega_{pi}^{-1}$$

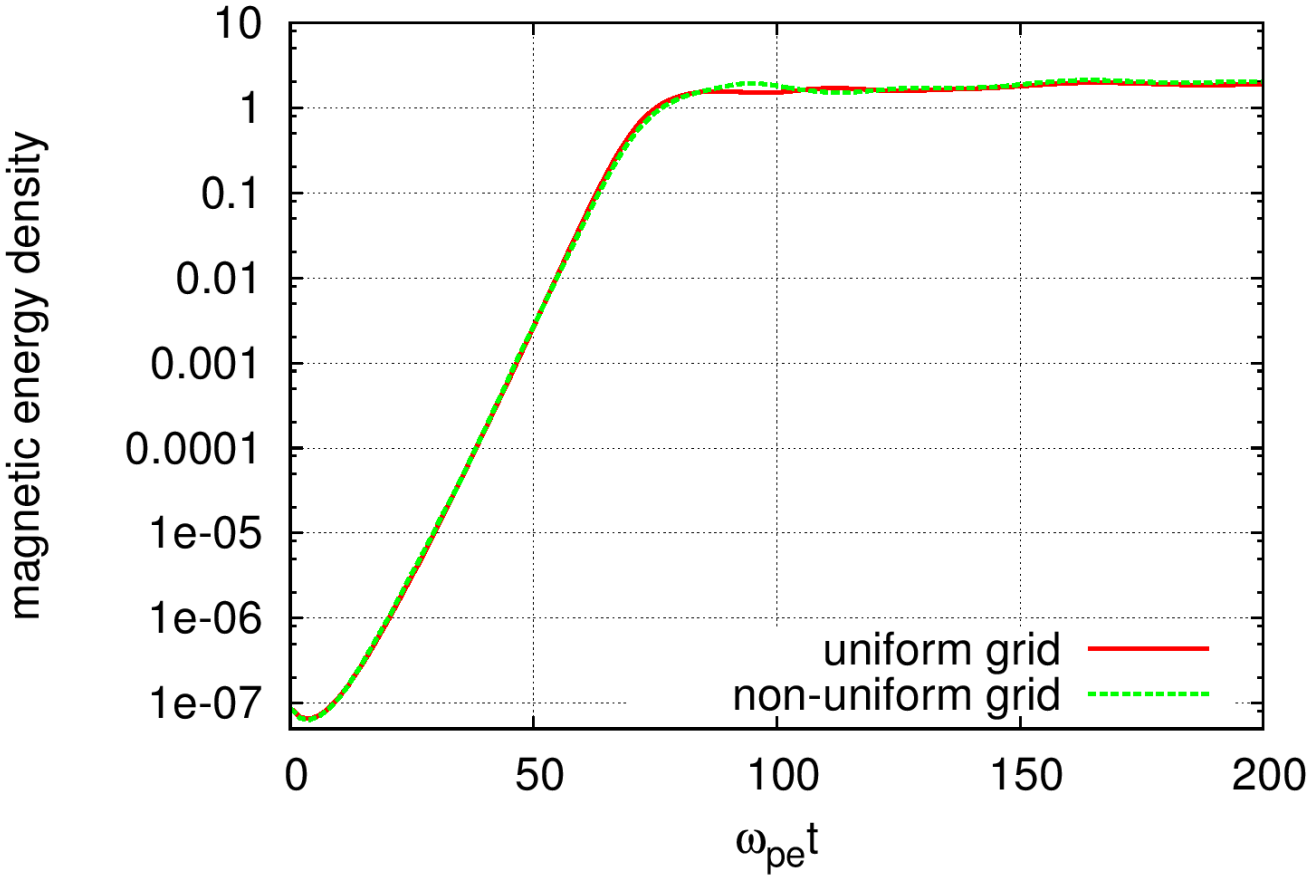

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

2D Electron Weibel instability: preconditioner performance

$$L_x \times L_y = 22 \times 22\ (d_e^2),\ N_{pc} = 200,\ \Delta t = 0.1\omega_{pi}^{-1}$$

$$N_x \times N_y = 128 \times 128$$

m_i/m_e	no preconditioner		with preconditioner	
	Newton	GMRES	Newton	GMRES
25	5.8	192.5	3	0
100	5.7	188.8	3	0
1836	7.7	237.8	4	2.8

$$m_i/m_e = 1836$$

$N_x \times N_y$	no preconditioner		with preconditioner	
	Newton	GMRES	Newton	GMRES
16×16	3.7	20	3	0.9
32×32	4	38.5	3	0.9
64×64	4.3	79.9	3	0.2

Luis Chacon, chacon@lanl.gov

2D KAW: impact of magnetization

$$B_{x,0} = 0.0667,\ v_{eT}/c = 0.0745\ (\beta_e = 0.1),\ \Delta t = 0.1\omega_{pi}^{-1}$$

$$L_x \times L_y = 10 \times 10(d_i^2)$$
$$N_{pc} = 500$$
$$N_x \times N_y = 64 \times 64$$
$$m_i/m_e = 25$$

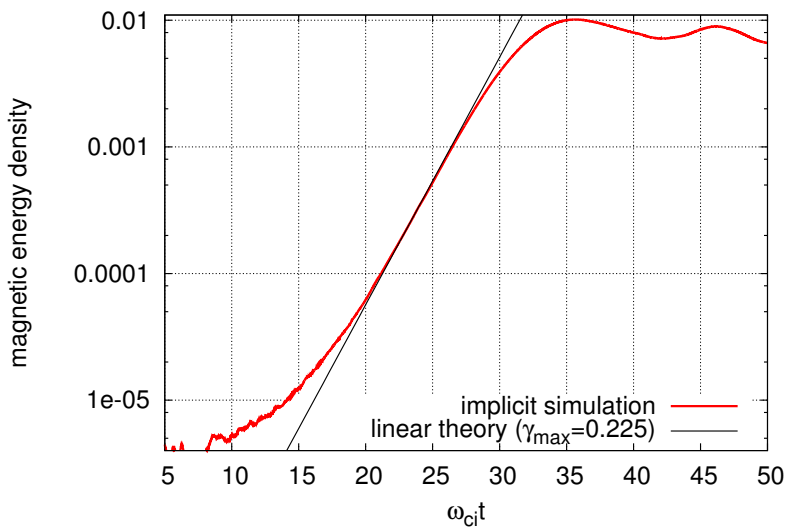

$$L_x \times L_y = 22 \times 22(d_i^2)$$
$$N_{pc} = 200,\ N_x \times N_y = 32 \times 32$$

Fixed magnetic field

m_i/m_e	no preconditioner		with preconditioner	
	Newton	GMRES	Newton	GMRES
25	4	171.9	3.2	1
150	4.5	764	4	2.9
600	7.4	4054.8	4	11.9

$$\omega_{pe}/\omega_{ce} = 3$$

m_i/m_e	no preconditioner		with preconditioner	
	Newton	GMRES	Newton	GMRES
150	4.5	738	4	3
600	5.8	1887	4	3.9
1836	NC	NC	4	5.9

Los Alamos NATIONAL LABORATORY EST.1943

Luis Chacon, chacon@lanl.gov

Summary and conclusions

- We have demonstrated a fully implicit, fully nonlinear, multidimensional PIC formulation that features:
  - **Exact local charge conservation** (via a novel particle mover strategy).
  - **Exact global energy conservation** (no particle self-heating or self-cooling).
  - **Adaptive particle orbit integrator** to control errors in momentum conservation.
  - **Canonical momenta** (EM-PIC only, reduced dimensionality).
- The approach is free of numerical instabilities: $\omega_{pe}\Delta t \gg 1$, and $\Delta x \gg \lambda_D$
  - Requires many fewer dofs (vs. explicit PIC) for comparable accuracy in challenging problems
  - Significant CPU gains (vs explicit PIC) have been demonstrated
  - The method has much potential for efficiency gains vs. explicit in long-time-scale applications:

$$\frac{CPU_{ex}}{CPU_{imp}} \sim \underbrace{\frac{1}{(k\lambda_D)^d}\frac{c}{v_{th,e}}}_{\text{Physics}} \overbrace{\frac{1}{N_{FE}}}^{\text{Precond.}} .$$

  - Moment-based acceleration is effective in minimizing N_{FE}, leading to an optimal algorithm.

Luis Chacon, chacon@lanl.gov