Speeding up simulations of relativistic systems using an optimal boosted frame

J.-L. Vay[1,3], W. M. Fawley[1], C. G. R. Geddes[1], E. Cormier-Michel[1], D. P. Grote[2,3]

[1]Lawrence Berkeley National Laboratory, CA

[2]Lawrence Livermore National Laboratory, CA

[3]Heavy Ion Fusion Science Virtual National Laboratory

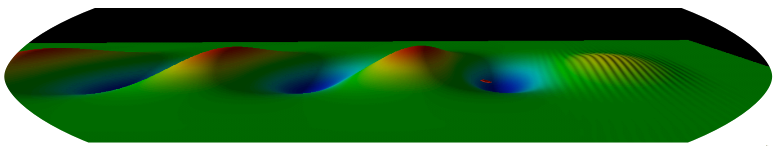

LWFA

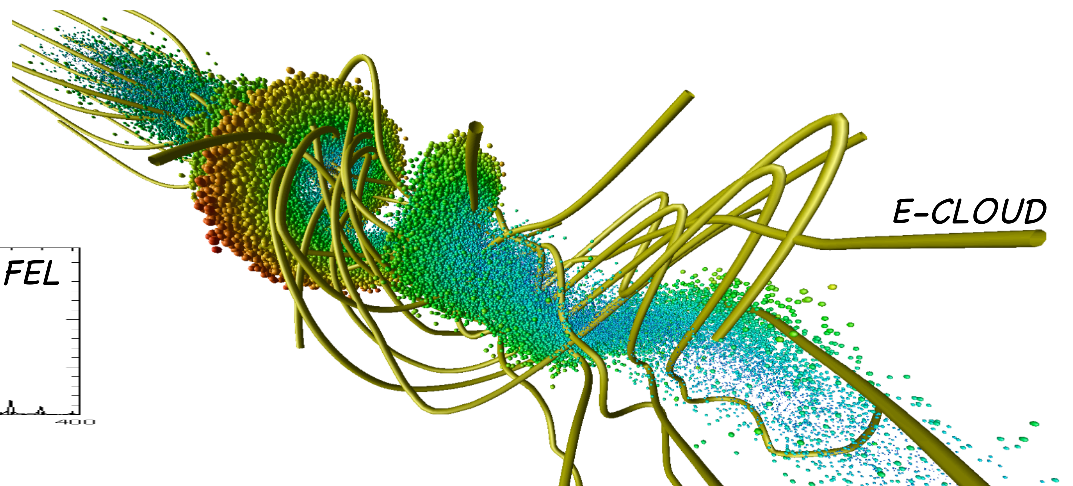

E-CLOUD

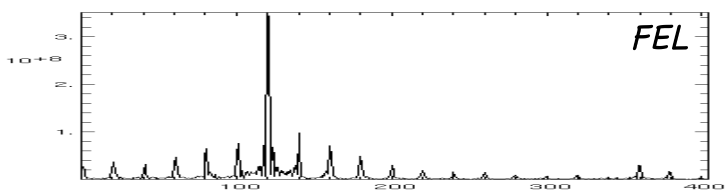

Outline

- Concept and history
- Difficulties
- Examples of application
  - electron cloud effects
  - laser wakefield acceleration
  - free electron laser
- Recent history and conclusion

At ICAP'06, we presented an analysis of the cost of self-consistent simulations of e-cloud instability,

*Vay et al, ICAP'06

Can 3-D self-consistent compete with quasi-static mode?
- computational cost of full 3-D run in two frames -

Lab frame

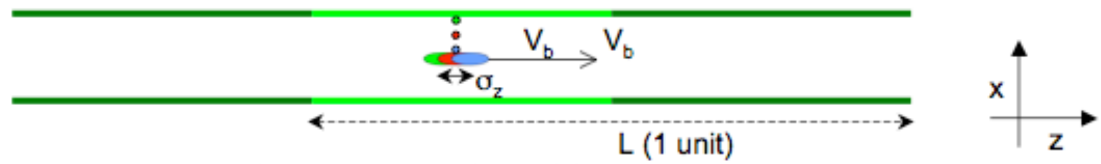

$\delta x = \sigma_x/n$; $\delta z = \min(\sigma_z, L)/n$

$\delta t < \min[\ \delta x/\max(v_x), \delta z/\max(v_z)\ ]$;

$T_{max} = N_{units} \times L/V_b$

$N_{op} = N_e \times T_{max}/\delta t$

Frame γ

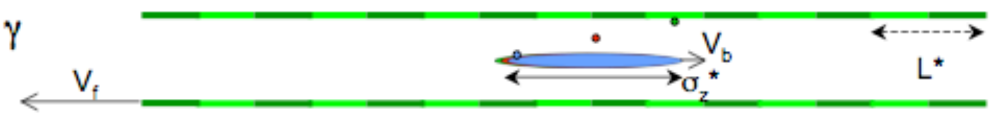

$\delta x^* = \sigma_x/n$; $\delta z^* = \min(\sigma_z^*, L^*) = \gamma\delta z$

$\delta t^* < \min[\ \delta x^*/\max(v_x^*), \delta z^*/\max(v_z^*)\ ] = \min[\ \delta x/(\max(v_x/\gamma)), \gamma\delta z/v_z] = \gamma\delta t$

$T^*_{max} = N_{units} \times L^*/(V_b - V_f) \sim T_{max}/\gamma$

$N^*_{op} = N_e \times T^*_{max}/\delta t^* \sim N_{op}/\gamma^2$

=> Computational cost greatly reduced in frame γ

contrasted it to the cost of quasistatic methods,

*Vay et al, ICAP'06

Comparison between quasi-static and full 3-D costs.

Lab frame

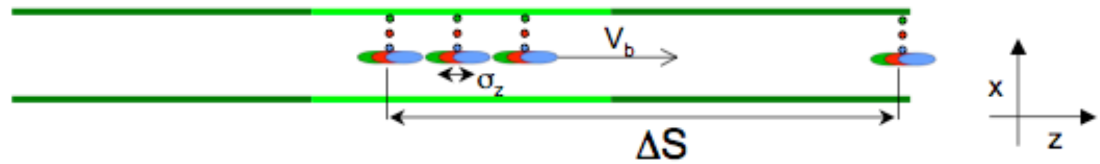

Quasi-static (HEADTAIL, QUICKPIC):

$\alpha \sim \Delta S/\sigma_z$

$N_{op,qs} = N_{op}/\alpha$

Frame γ

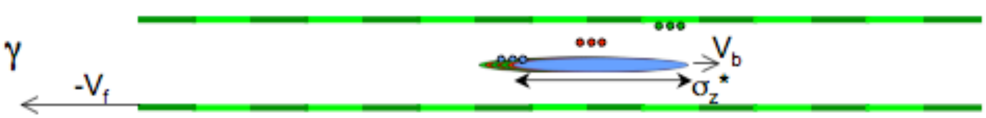

if $\sigma_z^* = \Delta S^*$, $\gamma^2 = \alpha$, $N^*_{op} = N_{op,qs}$

=> cost of full 3-D run in frame γ = cost of quasi-static mode in lab frame

and indicated broader applicability.

*Vay et al, ICAP'06

Conclusion

- We developed a unique combination of tools to study ECE
- WARP/POSINST code suite
  - Parallel 3-D PIC-AMR code with accelerator lattice follows beam self-consistently with gas/electrons generation and evolution,
- HCX experiment adresses ECE fundamentals (HIF/HEDP/HEP)
  - highly instrumented section dedicated to e-cloud studies,
  - extensive methodical benchmarking of WARP/POSINST,
- Being applied outside HIF/HEDP, to HEP accelerators
  - LHC, Fermilab MI, ILC,
  - Implemented "quasi-static" mode for direct comparison to HEADTAIL/QUICKPIC,
  - fund that self-consistent calculation has similar cost than quasi-static mode if done in moving frame (with γ>>1), thanks to relativistic contraction/dilatation bridging space/time scales disparities (applies to FEL, laser-plasma acceleration, plasma lens,…).

Range of space and time scales spanned by two identical beams crossing each other

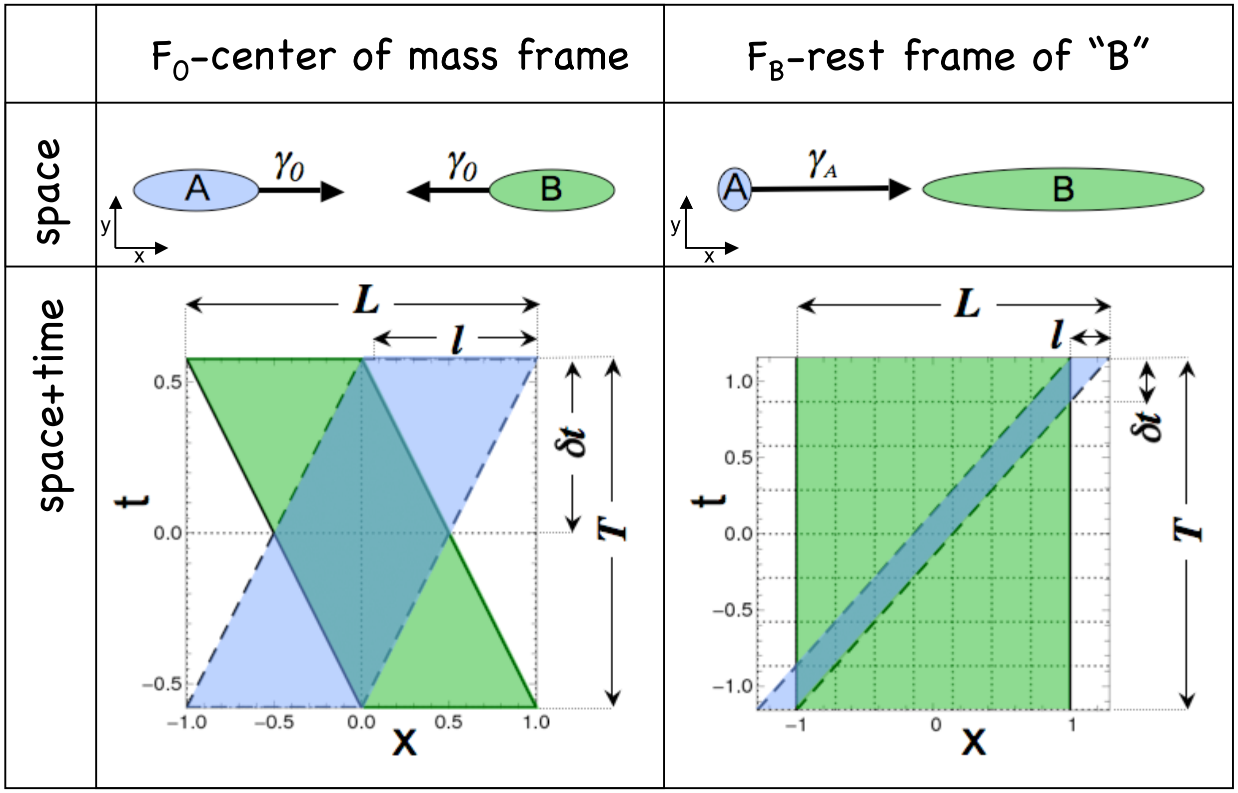

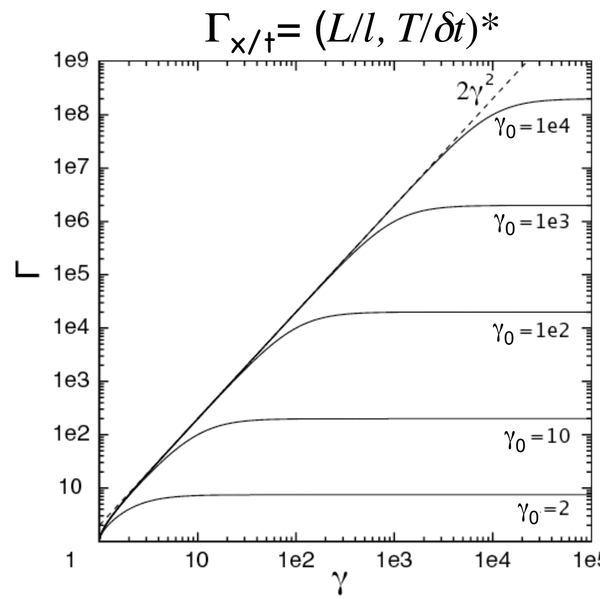

- Γ is not invariant under the Lorentz transformation: $\Gamma_{x/t} \propto \gamma^2$.
- There exists an "optimum" frame which minimizes it.
- Result is general and applies to light beams too.

*J.-L. Vay, *Phys. Rev. Lett.* **98**, 130405 (2007)

Consequence for computer simulations

of computational steps grows with the full range of space and time scales involved

Choosing optimum frame of reference to minimize range can lead to dramatic speed-up for relativistic matter-matter or light-matter interactions.

Calculation of e-cloud induced instability of a proton bunch*

- Proton energy: γ=500 in Lab
- L=5 km, continuous focusing

Code: Warp (Particle-In-Cell)

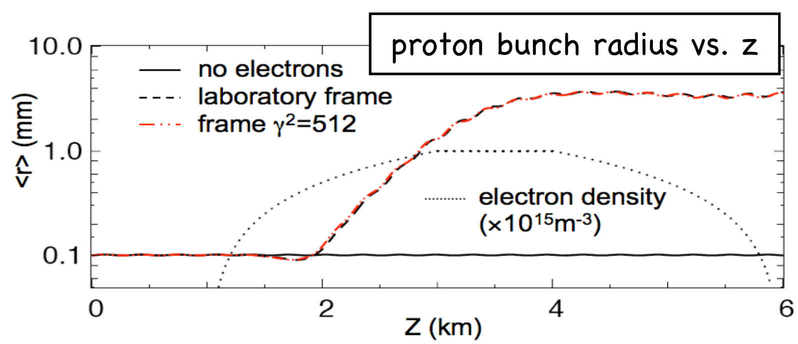

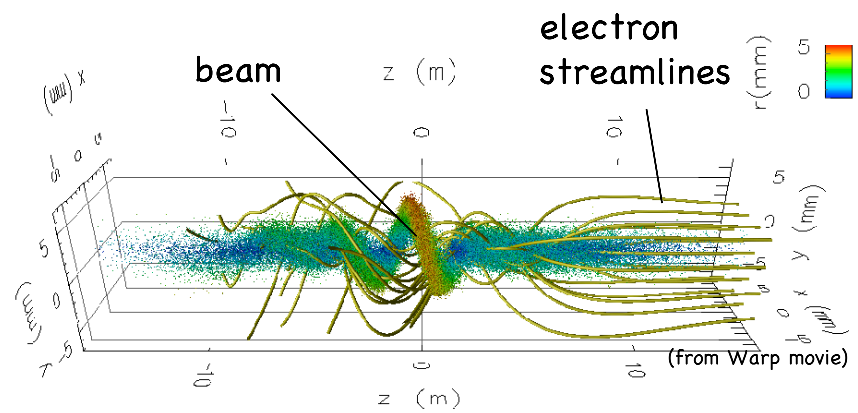

CPU time (2 quad-core procs):

- lab frame: >2 weeks
- frame with γ^2=512: <30 min

Speedup x1000

*J.-L. Vay, *Phys. Rev. Lett.* **98**, 130405 (2007)

Early history

- **1994**: Mori et al (UCLA) test 1D simulations (code Wake) of LWFA/PWFA in boosted frame but discontinue due to instability attributed to unresolved backward radiation (unpublished)

- **1999**: Ng (SLAC) & Vay (LBNL) perform 3D simulations (code BPIC) of beam focused by a plasma lens in boosted (beam) frame, estimated speedup >x10 (unpublished)

- **2006**:
  - Vay & Fawley (LBNL) perform 2D simulation (code Warp) of FEL toy problem in boosted (bucket) frame, estimated speedup ~x45,000 (progress report SBIR)
  - ICAP 06: Vay (LBNL), Friedman & Grote (LLNL) discuss calculation of e-cloud in boosted frame and contrast with quasistatic speedup, mention application to LWFA, FEL, plasma lenses,… (Proc. ICAP 2006)

- **2007**:
  - derivation of scaling showing γ^2 dependency for generic, e-cloud, LWFA and FEL problems; 1,000x demonstrated speedup on 3D simulation (code Warp) of e-cloud driven instability (Vay *PRL* **98**, 130405)
  - novel particle and field solver for boosted e-cloud simulations (Vay *PoP* **15**, 056701)
  - "passionate discussions" concerning feasability of application to LWFA because of upshifted backward radiation in boosted frame

Outline

- Concept and history
- Difficulties
- Examples of application
  - electron cloud effects
  - laser wakefield acceleration
  - free electron laser
- Recent history and conclusion

Seems simple but ⚠. Algorithms which work in one frame may break in another. Example: the Boris particle pusher.

- Boris pusher ubiquitous
  - In first attempt of e-cloud calculation using the Boris pusher, the beam was lost in a few betatron periods!
  - Position push: $\mathbf{X}^{n+1/2} = \mathbf{X}^{n-1/2} + \mathbf{V}^n \Delta t$ -- no issue
  - Velocity push: $\gamma^{n+1}\mathbf{V}^{n+1} = \gamma^n\mathbf{V}^n + \frac{q\Delta t}{m}\left(\mathbf{E}^{n+1/2} + \frac{\gamma^{n+1}\mathbf{V}^{n+1} + \gamma^n\mathbf{V}^n}{2\,\gamma^{n+1/2}} \times \mathbf{B}^{n+1/2}\right)$

  issue: **E**+v×**B**=0 implies **E**=**B**=0 => large errors when **E**+v×**B**≈0 (e.g. relativistic beams).

- Solution
  - Velocity push: $\gamma^{n+1}\mathbf{V}^{n+1} = \gamma^n\mathbf{V}^n + \frac{q\Delta t}{m}\left(\mathbf{E}^{n+1/2} + \frac{\mathbf{V}^{n+1} + \mathbf{V}^n}{2} \times \mathbf{B}^{n+1/2}\right)$

- Not used before because of implicitness. We solved it analytically*

$$\begin{cases} \gamma^{i+1} = \sqrt{\dfrac{\sigma + \sqrt{\sigma^2 + 4(\tau^2 + u^{*2})}}{2}} \\ \mathbf{u}^{i+1} = [\mathbf{u}' + (\mathbf{u}' \cdot \mathbf{t})\mathbf{t} + \mathbf{u}' \times \mathbf{t}]/(1+t^2) \end{cases}$$

(with $\mathbf{u} = \gamma\mathbf{v}$, $\mathbf{u}' = \mathbf{u}^i + \frac{q\Delta t}{m}\left(\mathbf{E}^{i+1/2} + \frac{\mathbf{v}^i}{2} \times \mathbf{B}^{i+1/2}\right)$, $\boldsymbol{\tau} = (q\Delta t/2m)\mathbf{B}^{i+1/2}$, $u^* = \mathbf{u}' \cdot \boldsymbol{\tau}/c$, $\sigma = \gamma'^2 - \tau^2$, $\gamma' = \sqrt{1 + u'^2/c^2}$, $\mathbf{t} = \boldsymbol{\tau}/\gamma^{i+1}$).

*J.-L. Vay, *Phys. Plasmas* **15**, 056701 (2008)

Other possible complication: inputs/outputs

- Often, initial conditions known and output desired in laboratory frame
  - relativity of simultaneity => inject/collect at plane(s) ⊥ to direction of boost.
- Injection through a moving plane in boosted frame (fix in lab frame)
  - fields include frozen particles,
  - same for laser in EM calculations.

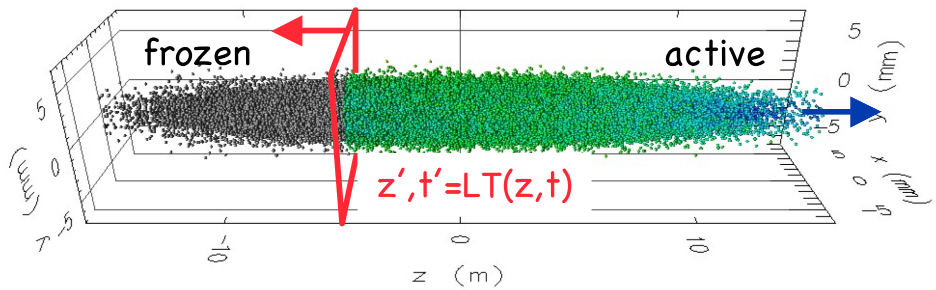

- Diagnostics: collect data at a collection of planes
  - fixed in lab fr., moving in boosted fr.,
  - interpolation in space and/or time,
  - already done routinely with Warp

  for comparison with experimental data, often known at given stations in lab.

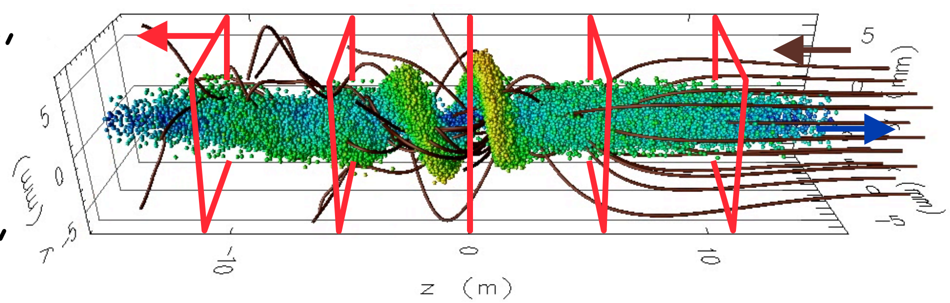

Outline

- Concept and history
- Difficulties
- Examples of application
  - electron cloud effects
  - laser wakefield acceleration
  - free electron laser
- Recent history and conclusion

E-cloud: benchmarking against quasistatic model for LHC scenario

The "quasistatic" approximation uses the separation of time scales for pushing beam and e-cloud macro-particles with different "time steps": used in QuickPIC (USC/UCLA), Headtail (CERN), PEHTS (KEK), CMAD (SLAC), Warp (LBNL), ...

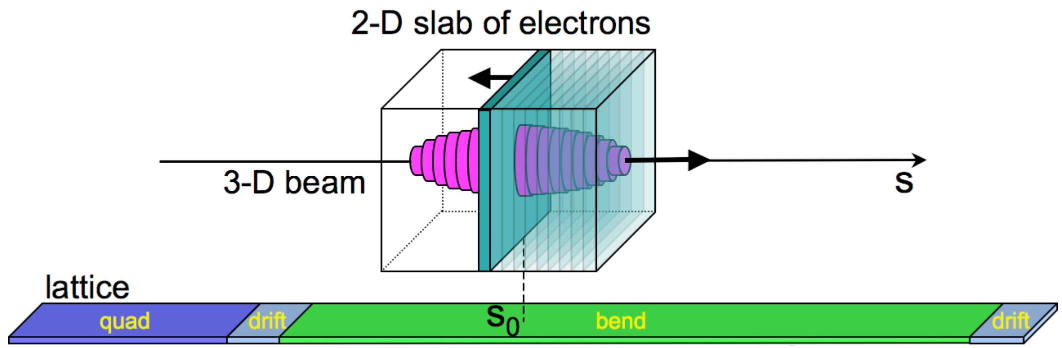

Excellent agreement on emittance growth between boosted frame full PIC and "quasistatic" for e-cloud driven transverse instability in continuous focusing model of LHC:

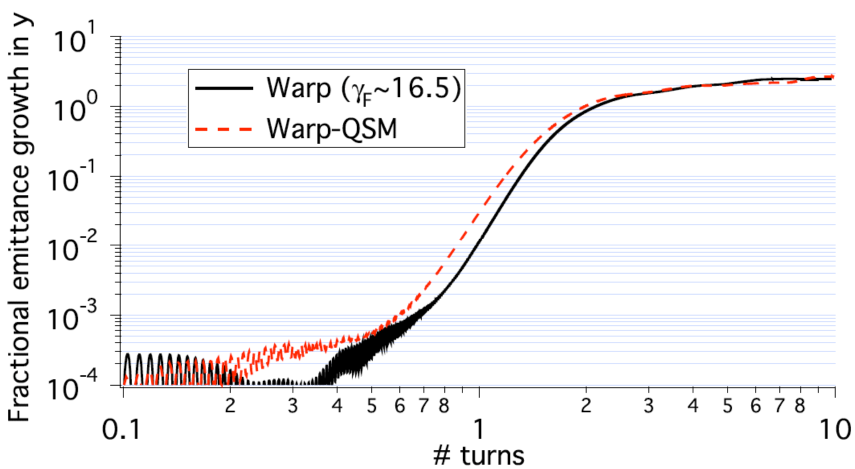

The 2 runs have similar computational cost, thus how to choose one method other another?

- boosted frame method offers less approximation to the physics, which may matter in some cases,
- parallelization of quasistatic codes more complicated due to pipelining in the longitudinal direction.

Outline

- Concept and history
- Difficulties
- Examples of application
  - electron cloud effects
  - laser wakefield acceleration
  - free electron laser
- Recent history and conclusion

BELLA 40 J PW Laser - Components for a Laser Plasma Collider

10 GeV stages

Energy spread &
Emittance preservation

Injection + Staging

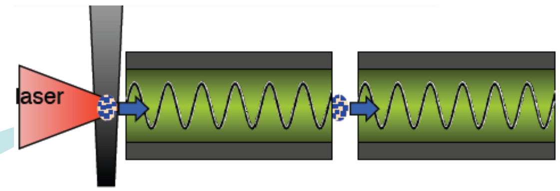

BELLA
PW laser
40 J / 40 fs

Positron acceleration + PWFA expt.'s

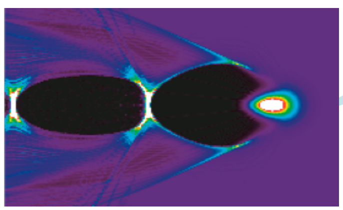

Radiation sources

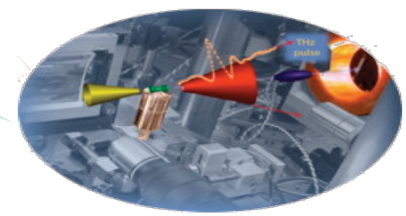

Simulating 10 GeV stages explicitly (PIC) in lab frame needs ~1G CPU•hours ⇒ impractical*

Predictions have relied on theory, reduced models (fluid, envelope, quasistatic), scaling:

- Energy gain $\propto n^{-1}$: 100 MeV at 10^{19}/cc ⇒ 10 GeV at 10^{17}/cc
- Length $\propto n^{-3/2}$: 1mm at 10^{19}/cc ⇒ 1m at 10^{17}/cc
- Gradient $\propto n^{1/2}$: 100 GV/m at 10^{19}/cc ⇒ 10 GV/m at 10^{17}/cc

Can simulations of full scale 10 GeV stages be practical using a Lorentz boosted ref. frame?

- difficulty: backward emitted radiation frequency upshifted in boosted frame, (noise, instabilities).

* Cormier-Michel et al, Proc. AAC 2008; Geddes et al, Proc. PAC'09

BERKELEY LAB

2D scaled simulations of a 10 GeV class LWFA stage
(λ=0.8μm, a_0=1, k_pL=2, L_p=1.5mm in lab)

Snapshots of surface plot of // electric field in lab frame and boosted frame at γ=10

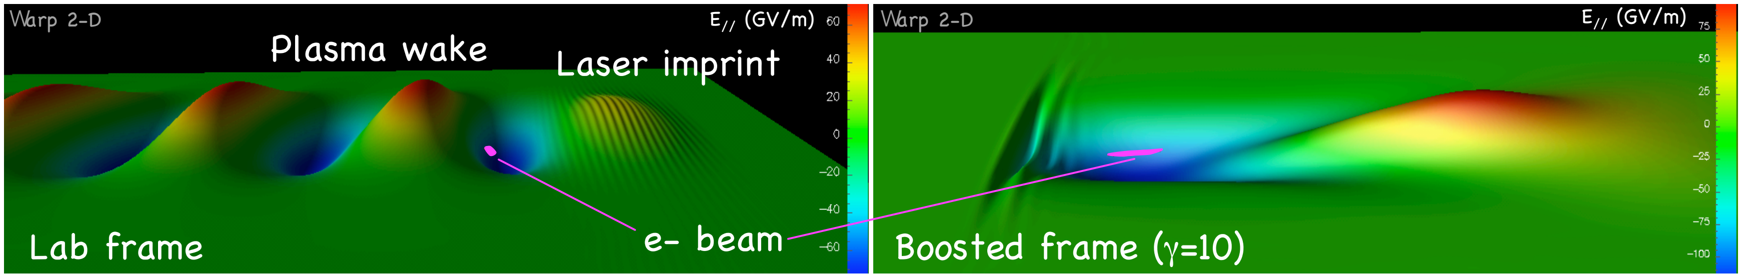

2D ⊥ and // electric field history in lab frame

Station z=0.154mm

Station z=1.354mm

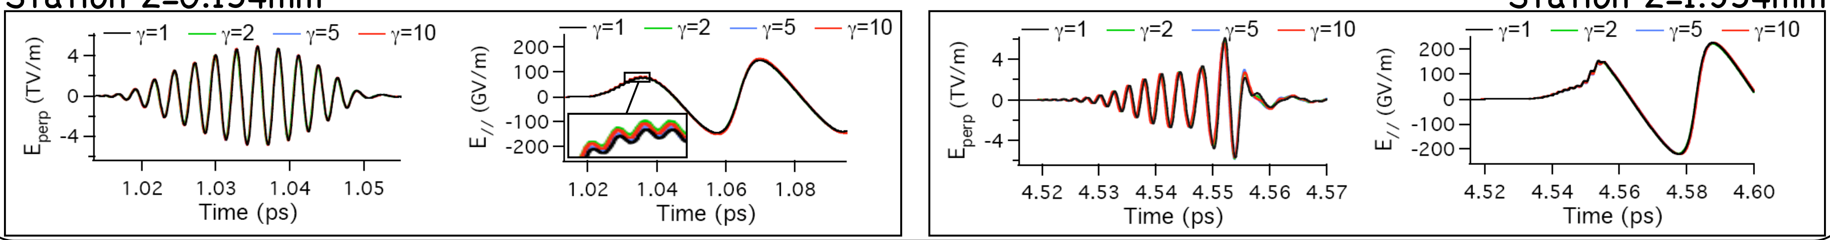

2D Average beam energy and CPU time vs position in lab frame

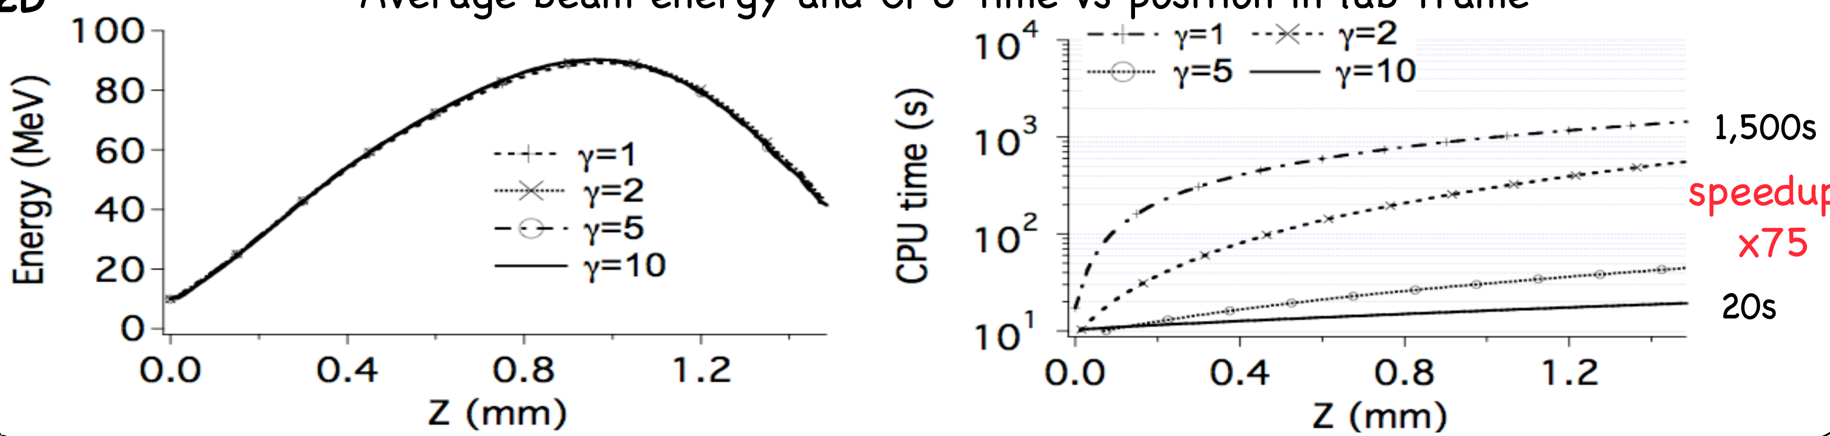

3D scaled simulations of a 10 GeV LWFA stage
(λ=0.8μm, a_0=1, k_pL=2, L_p=1.5mm in lab)

Snapshots of surface plot of // electric field in lab frame and boosted frame at γ=10

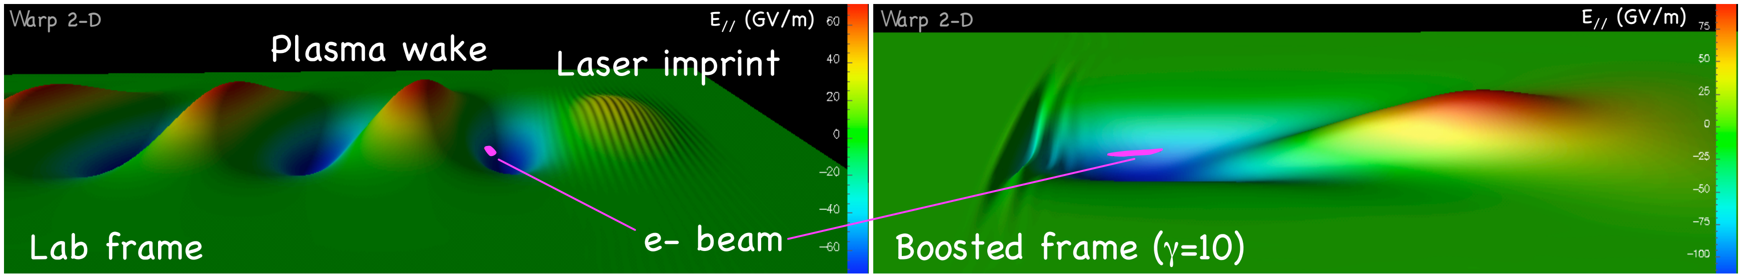

3D

⊥ and // electric field history in lab frame

Station z=0.3mm

Station z=1.354mm

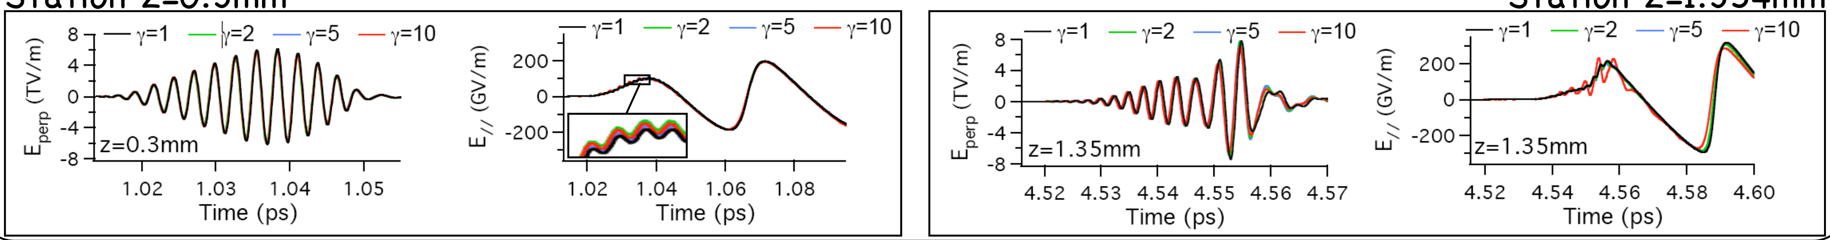

3D

Average beam energy and CPU time vs position in lab frame

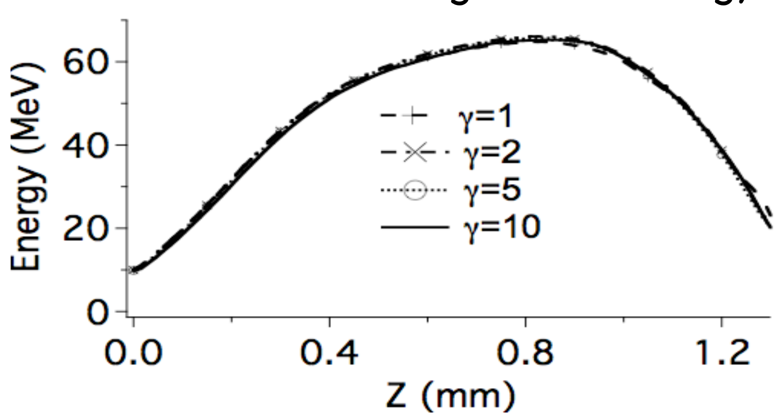

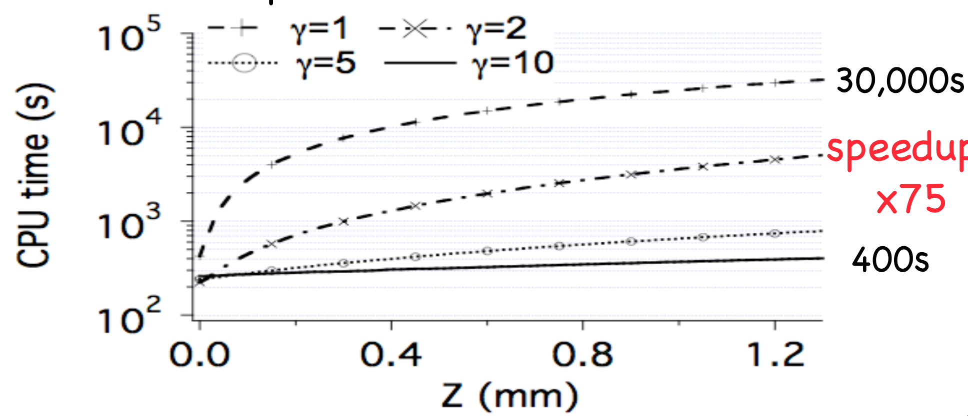

Full scale simulations of a 10 GeV LWFA stage

Simulations in 1D/2D/3D at plasma densities of 10^{19}cc, 10^{18}cc and 10^{17}cc show good agreement on (scaled) beam energy gain:

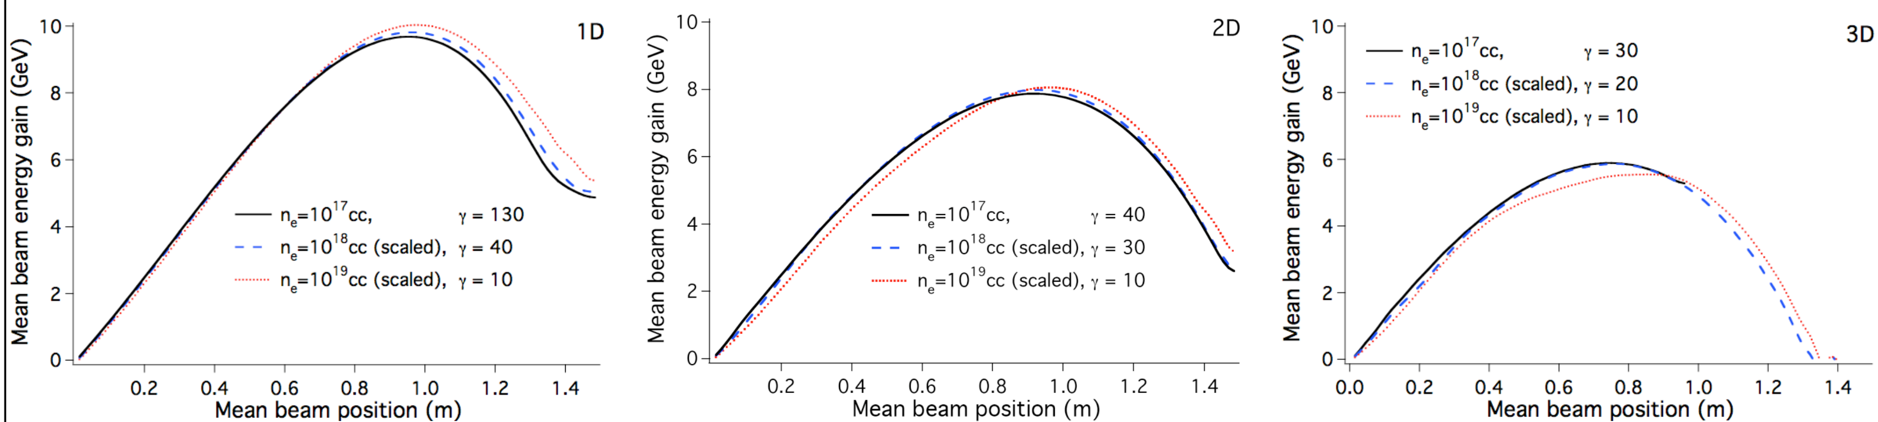

- 1D: max γ_{frame} = 130 ⇒ speedup > 10,000
- 2D: max γ_{frame} = 40 ⇒ speedup > 1,000
- 3D: max γ_{frame} = 30 ⇒ speedup > 500
  - 24h using 256 CPUs ⇒ more than one year x 256 CPUs in lab frame!

Max γ_{frame} achieved in 2D and 3D limited by instability developing at front of plasma

origin (numerical Cerenkov?) and cures are being studied.

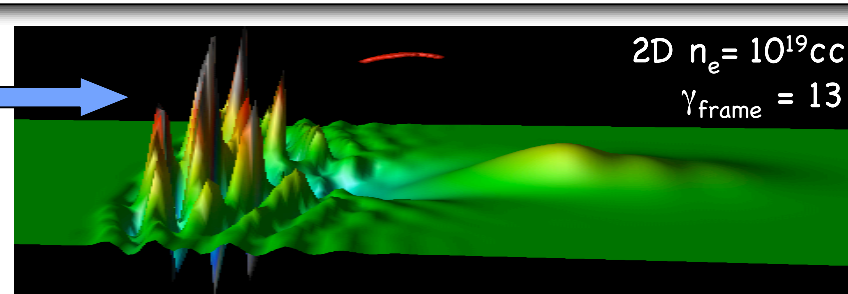

*similar work by Bruhwiler et al (Tech X), Martins et al (UCLA/IST)

Is the instability due to Yee solver numerical dispersion errors? Implementation of a low-dispersion solver in Warp*

Enlarged stencil** => no disp. in x,y,z

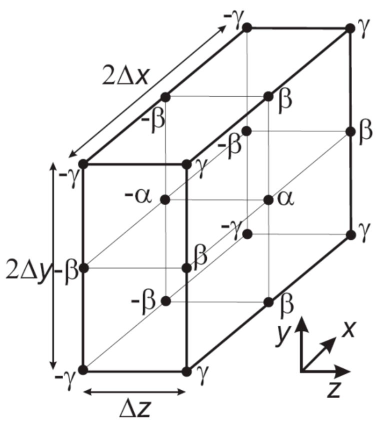

$$E_x|^{n+1}_{i+1/2,j,k} = E_x|^{n}_{i+1/2,j,k} - \alpha\frac{\Delta t}{\epsilon_0}D_{z,0}H_y|^{n+1/2}_{i+1/2,j,k} - 4\beta\frac{\Delta t}{\epsilon_0}D_{z,1}H_y|_{i+1/2,j,k} - 4\gamma\frac{\Delta t}{\epsilon_0}D_{z,2}H_y|^{n+1/2}_{i+1/2,j,k} + \alpha\frac{\Delta t}{\epsilon_0}D_{y,0}H_z|^{n+1/2}_{i+1/2,j,k} + 4\beta\frac{\Delta t}{\epsilon_0}D_{y,1}H_z|^{n+1/2}_{i+1/2,j,k} + 4\gamma\frac{\Delta t}{\epsilon_0}D_{y,2}H_z|^{n+1/2}_{i+1/2,j,k}.$$

Stencil on H unchanged

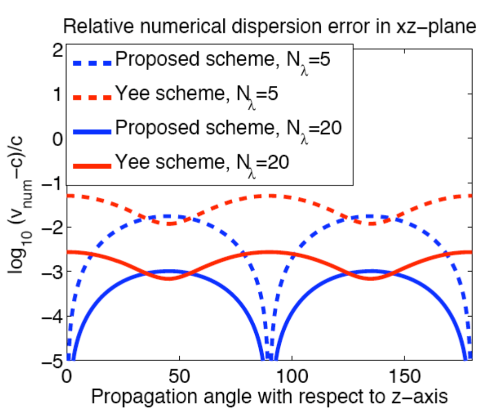

Issues for PIC:

- source terms are not given,
- odd-even oscillation when $\Delta t=\Delta x/c$.

Implementation in Warp

$$H_x|^{n+1}_{i+1/2,j,k} = H_x|^{n}_{i+1/2,j,k} - \alpha\frac{\Delta t}{\epsilon_0}D_{z,0}E_y|^{n+1/2}_{i+1/2,j,k} - 4\beta\frac{\Delta t}{\epsilon_0}D_{z,1}E_y|_{i+1/2,j,k} - 4\gamma\frac{\Delta t}{\epsilon_0}D_{z,2}E_y|^{n+1/2}_{i+1/2,j,k} + \alpha\frac{\Delta t}{\epsilon_0}D_{y,0}E_z|^{n+1/2}_{i+1/2,j,k} + 4\beta\frac{\Delta t}{\epsilon_0}D_{y,1}E_z|^{n+1/2}_{i+1/2,j,k} + 4\gamma\frac{\Delta t}{\epsilon_0}D_{y,2}E_z|^{n+1/2}_{i+1/2,j,k}.$$

- E and H switched,
  => E push same as Yee,
- exact charge conservation preserved in 2D & 3D with unmodified Esirkepov current deposition and implied enlarged stencil on div E.

1-2-1 filter needed at Δt=Δx/c

1-D simulation of current Heaviside step

full amplitude odd-even oscillations for Δt=Δx/c

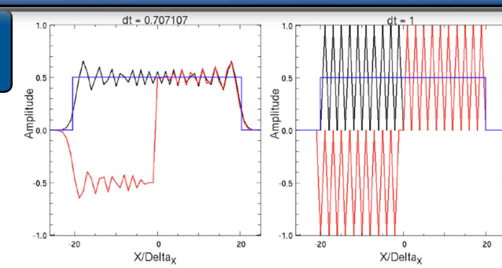

1-D simulation of LWFA at Δt=Δx/c

No filter

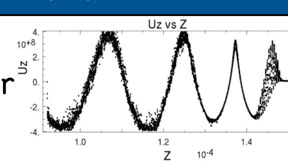

121 filter

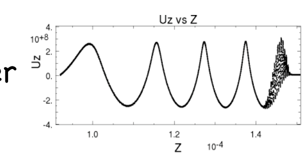

*Vay et al, poster on AMR-PIC, ICAP'09, Thursday

**M. Karkkainen, et al., Proceedings of ICAP'06, Chamonix, France

Low dispersion solver reduces instability growth but is not sufficient

Longitudinal electric field from 2D simulations at plasma densities of 10^{18}cc with $\gamma_{frame} = 40$

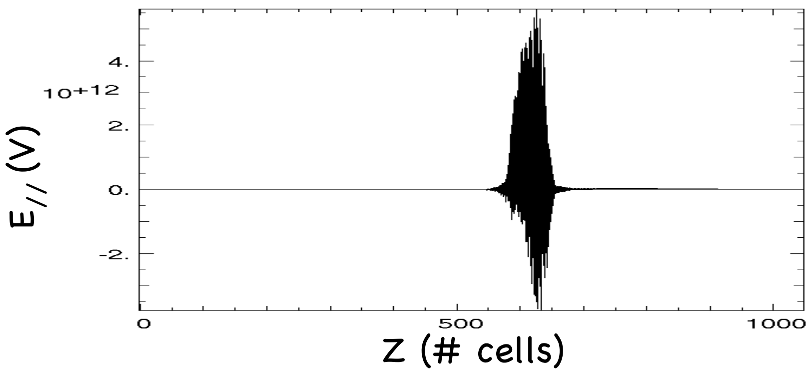

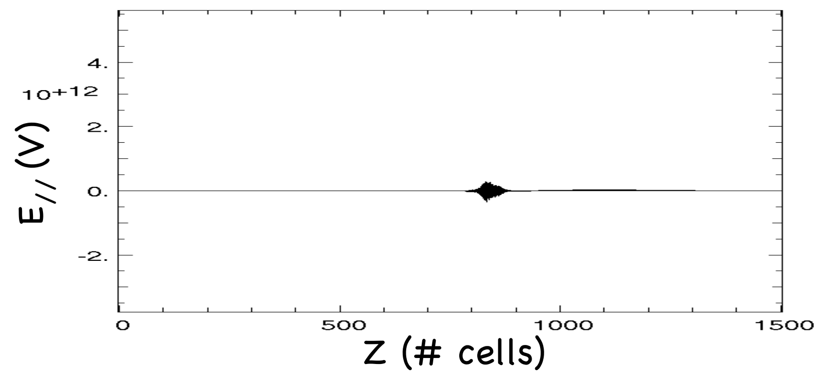

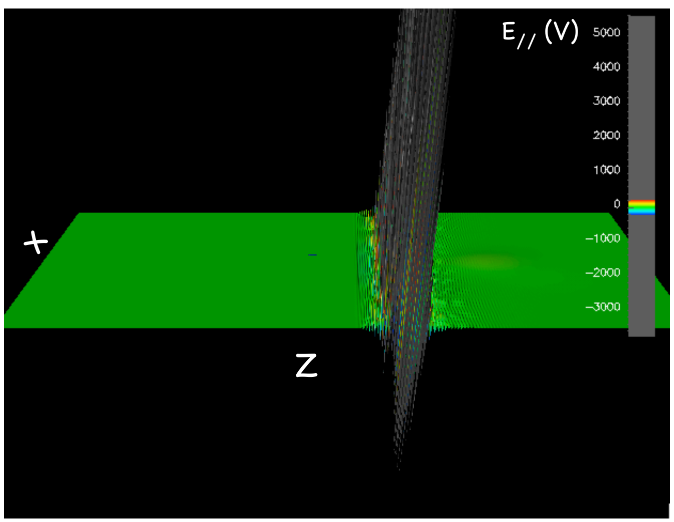

Yee field solver

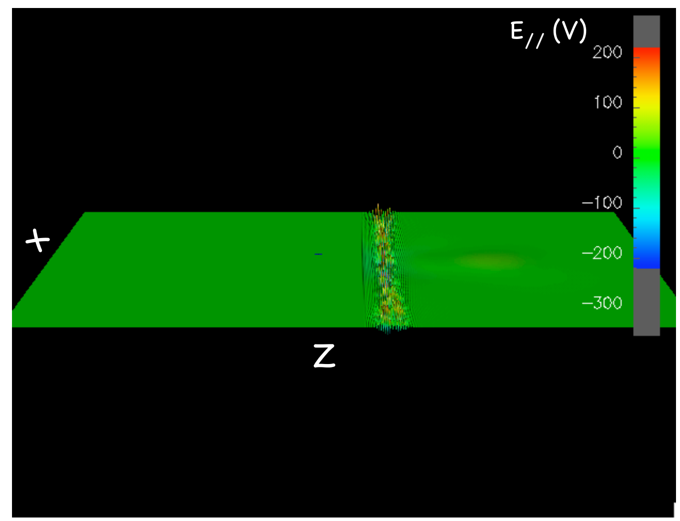

Karkkainen field solver

Outline

- Concept and history
- Difficulties
- Examples of application
    - electron cloud effects
    - laser wakefield acceleration
    - free electron laser
- Recent history and conclusion

FEL in Boosted-Frame E&M Code

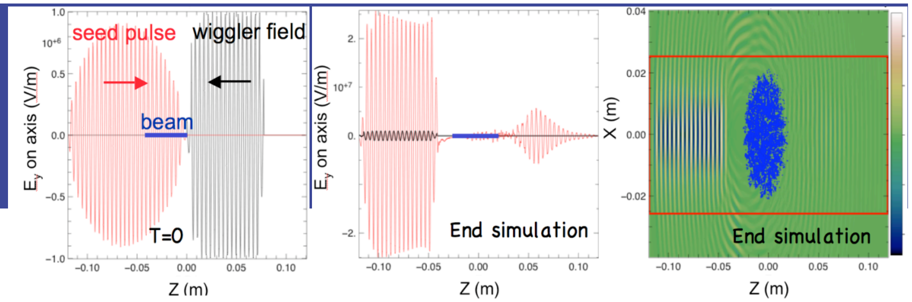

- Standard FEL codes use slowly-varying envelope (Eikonal) and wiggler-period averaring approximations; Physics ignored by Eikonal codes but accessible to boosted frame approach
  - Backward wave emission
  - Wide-angle emission (generally highly red-shifted)
  - CSE for all undulator, e-beam configurations (very short beams; beams with rapidly-varying envelope properties, beams bunched with "multiple colors")
  - Properties of "very" high gain systems ($L_G/\lambda_u < 5$)
  - FEL emission from beams in multiple harmonic undulators (biharmonic or triharmonic undulators; effects of adiabatic match sections)
  - FEL emission in waveguides where v_{group} strongly varying with ω (normally relevant to microwave FEL's operating near cutoff)
- Benchmarking of BF with Eikonal code Ginger*: impressive speedup compared to *full* E&M *but much slower than standard eikonal method*
  - Not likely to become dominant paradigm for short wavelength FEL's but *might* be useful for very high gain microwave/far-IR devices or situations with wideband spectral output

*W.M. Fawley et al, Proc AAC'08; Proc. PAC'09

Outline

- Concept and history
- Difficulties
- Examples of application
  - electron cloud effects
  - laser wakefield acceleration
  - free electron laser
- Recent history and conclusion

Recent history

- **2007:**
  - Martins (IST) et al: 3D simulations (code Osiris) of 10 GeV LWFA stage with speedup ~x300 (Martins et al, APS-DPP)

- **2008:**
  - Bruhwiler et al (Tech X): 1D/2D simulations (code Vorpal) of 10 GeV LWFA stage with speedup ~x2,000 (Bruhwiler et al, Proc AAC'08)
  - Martins (IST et al): 3D simulations (code Osiris) of 10 GeV LWFA stage with speedup ~x300 (Martins et al, Proc AAC'08)
  - Fawley et al (LBNL): 2D simulations (code Warp) of FEL with detailed benchmarking against Eikonal code Ginger (Fawley et al, Proc AAC'08)

- **2009:**
  - Martins (IST) et al: 3D simulations (code Osiris) of 25+ GeV LWFA stage with speedup ~x300 (Martins et al, Proc PAC'09)
  - Vay et al (LBNL): 1D/2D/3D simulations (code Warp) of 10 GeV LWFA stage with speedup up to ~x1,000 (2D/3D) and >x10,000 (1D) (Vay et al, Proc PAC'09; Proc Scidac'09, slides APS-DPF)
  - Fawley et al (LBNL): 3D simulations (code Warp) of coherent synchrotron radiation (CSR) (Fawley et al, unpublished)

Conclusion and outlook

- The range of scales of a system is not a Lorentz invariant ($\propto\gamma^2$), and there exists an optimum frame minimizing it => orders of magnitude speedup predicted for some simulations.

- Calculating in a boosted frame more demanding, eventually:
  - developed new particle pusher for e-cloud problems,
  - added capabilities for injection/diagnostics in boosted frame.

- Orders of magnitude speedup demonstrated for a class of first-principle simulations of multiscale problems: laser-plasma acceleration, e-cloud in HEP accelerators, free electron lasers.

- Explore other applications: CSR, astrophysics,…

- Can we develop methods which costs do not depend on frame?