Scalable Arbitrary-Order Pseudo-Spectral Electromagnetic Solver

J.-L. Vay[1], T. Drummond[1], A. Koniges[1], B. B. Godfrey[1,2], I. Haber[2]
[1]Lawrence Berkeley National Laboratory, [2]University of Maryland

SUMMARY: While pseudo-spectral methods have been popular in the early PIC codes, the finite-difference time-domain method has become dominant with the rise of massively parallel computing owing to its locality advantage that lends to message passing that is limited to neighboring processors. Recently, a novel parallelization strategy was proposed [1] that takes advantage of the local nature of Maxwell equations that has the potential to combine pseudo-spectral accuracy with finite-difference favorable parallel scaling. In this talk, we will present the latest developments in the implementation of spectral-based solvers in Warp and discuss our latest findings.

Context

Particle-In-Cell

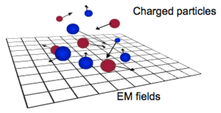

Method of choice for modeling of plasmas:

- based on first principles:
  → includes nonlinear, 3D, kinetic effects,
- particle push and EM solver are local:
  → scales well to >100ks cores,
- subject to instabilities (talk by B. Godfrey):
  → analyze and revisit methods.

Electromagnetic solvers

Finite-Difference Time Domain Warp (LBNL/LLNL/U. Maryland), Osiris (UCLA), V-sim (Tech-X), etc.	Pseudo-Spectral Time Domain UPIC-EMMA (UCLA)	Pseudo-Spectral Analytic Time Domain[1,2] Warp (LBNL/LLNL/U. Maryland)
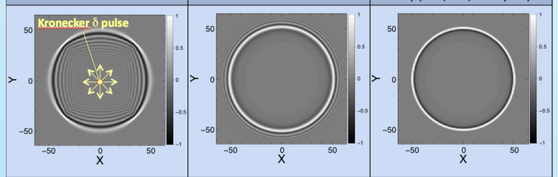		
• Numerical dispersion, • anisotropy, • Courant condition: $c\Delta t \leq 1/\sqrt{\frac{1}{\Delta x^2}+\frac{1}{\Delta y^2}+\frac{1}{\Delta z^2}}$	• Numerical dispersion, • isotropy, • Courant condition: $c\Delta t \leq 2/\pi\sqrt{\frac{1}{\Delta x^2}+\frac{1}{\Delta y^2}+\frac{1}{\Delta z^2}}$	• Exact dispersion, • isotropy, • Courant condition: **None**

But spectral solvers involve global operations:
→ harder to scale to large # of cores

Finite-Difference Time-Domain (FDTD)	Pseudo-Spectral Time-Domain (PSTD)
$B_z^{n+1} = B_z^n + \Delta t\left(\frac{\Delta E_x}{\Delta y} - \frac{\Delta E_y}{\Delta x}\right)$	$B_z^{n+1} = B_z^n + \mathcal{F}^{-1}\left(ik_y \mathcal{F}(E_x)\right) - \mathcal{F}^{-1}\left(ik_x \mathcal{F}(E_y)\right)$

$\mathcal{F}$=FFT

Finite Difference (FDTD)

local "cheap" communications

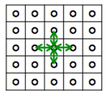

Easier to scale

Spectral

global "costly" communications

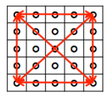

Harder to scale

New concept* opens the way to spectral accuracy with FDTD scaling: finite speed of light → local FFTs.

*J.-L. Vay, I. Haber, B. Godfrey, *J. Comput. Phys.* **243**, 260-268 (2013)

Novel Parallelization Concept

Explanation on single Kronecker pulse

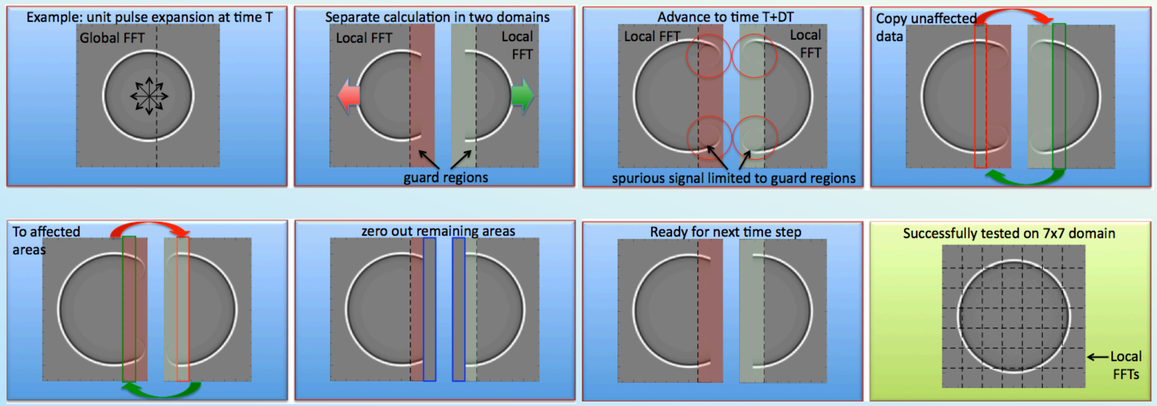

*J.-L. Vay, I. Haber, B. Godfrey, *J. Comput. Phys.* **243**, 260-268 (2013)

Arbitrary Order FDTD+PSTD Enables Tunable Solver

FDTD converges to PSTD when order goes to infinity and PSATD when time step goes to zero.

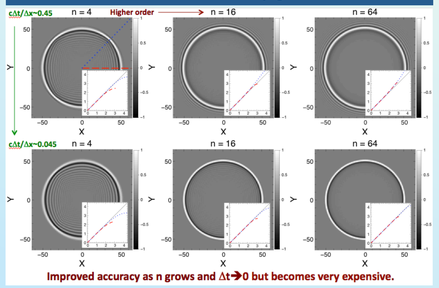

Improved accuracy as n grows and Δt→0 but becomes very expensive.

FDTD converges to PSTD when order goes to infinity and PSATD when time step goes to zero.

Step 1: multiply wave numbers by function to emulate FD stencil at any order

– k*=g(k,n)×k

$$B_z^{n+1} = B_z^n + \mathcal{F}^{-1}\left(ik_y^* \mathcal{F}(E_x)\right) - \mathcal{F}^{-1}\left(ik_x^* \mathcal{F}(E_y)\right)$$

Step 2: rewrite equations in matrix form to enable arbitrary time step via pre-computed subcycling

$$\begin{bmatrix} B_z \\ E_x \\ E_y \end{bmatrix}^{n+N} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix}^N \times \begin{bmatrix} B_z \\ E_x \\ E_y \end{bmatrix}^n$$

Ultimate flexibility is provided by runtime auto-tuning of order of accuracy vs locality and multi-level concurrency for a given architecture

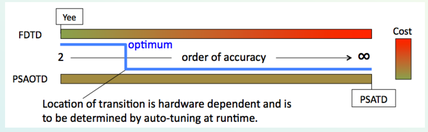

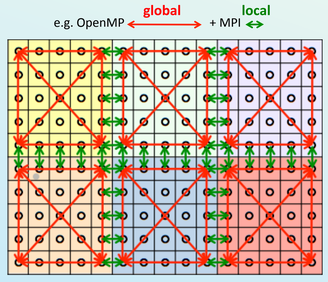

Flexible order of accuracy and subcycling of time steps enable control of locality

PSAOTD equivalent to FDTD at low order to machine precision

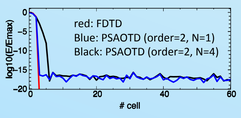

PSAOTD enables extreme accuracy in space and time while preserving locality

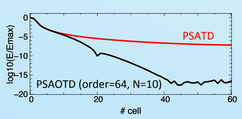

Scaling Study

Strong scaling of Warp's spectral PIC solver on a test problem of a 4,096x4,096 2-D plasma with periodic boundary conditions and 64 macroparticles/cell

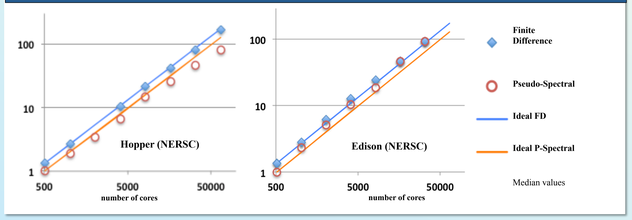

Performance variability for different runs of the solvers in Warp on Hopper and Edison. Performance plots are relative to the best runtime per configuration.

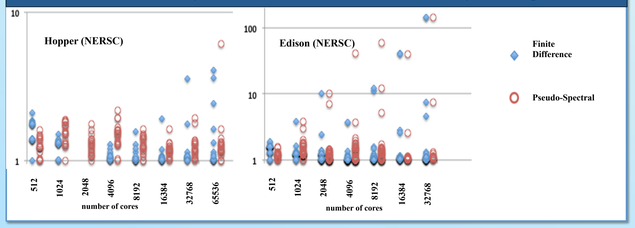

Strong scaling of the Parallel Pseudo-Spectral Solver based on domain decompositions vs. A global domain FFT on NERSC's Hopper

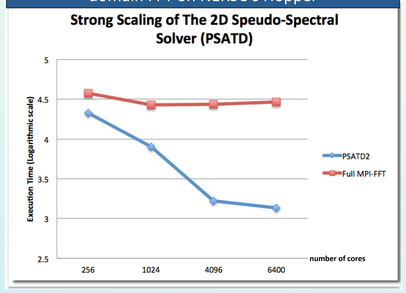

Hopper is a Cray XE6 with 6384 nodes. The nodes have two twelve-core AMD 'MagnyCours' 2.1-GHz processors per node, 24 cores per node (153,216 total cores) Peak Gflop/s rate is 8.4 Gflops/core. A single given compute node is always allocated to run a single user job; multiple jobs never share a compute node.

Edison is a Cray XC30 supercomputer with 5,576 computes nodes, 133,824 cores in total. It has 5576 nodes. Each node has two sockets, each socket is populated with a 12-core Intel "Ivy Bridge" processor, or 24 cores per node.

U.S. DEPARTMENT OF ENERGY | Office of Science
NERSC
ComPASS SciDAC-3