On elliptic solves for 3D XMHD

L. Chacón

Los Alamos National Laboratory
P.O. Box 1663, Los Alamos, NM 87545

CEMM MEETING
OCT. 29TH, 2006
PHILADELPHIA, PA USA

Reports of interest

1. Simulation-Based Engineering Science (SBES): Revolutionizing Engineering Science through Simulation, Feb. 2006, NSF Report
2. Fusion Simulation Project: Integrated Simulation & Optimization of Fusion Systems, Dec. 01, FESAC ISOFS Subcommittee Report

Outline

- Motivation
- Parabolization: key for SCALABILITY
- Application: Resistive 3D MHD
- Application: Hall MHD
- Implicit + AMR: proof of principle

Challenges in fusion simulation: “The tyranny of scales”

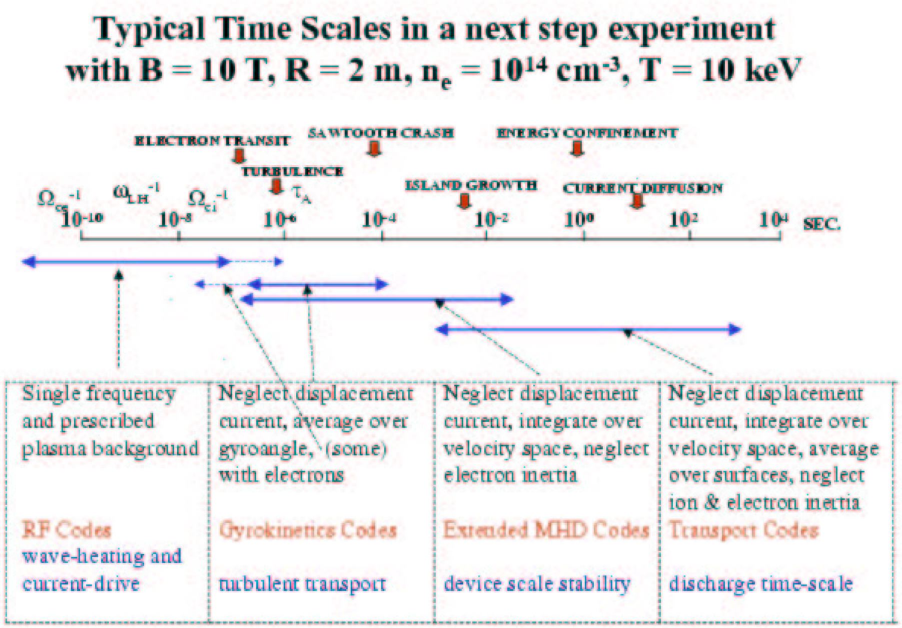

(a) Time scales in fusion plasmas (FSP report)

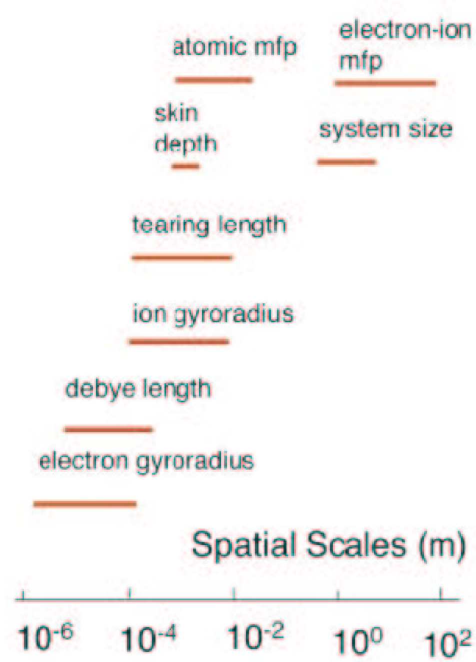

(b) Length scales in a typical fusion plasma (Tang, *Phys. Plasmas*, 9 (5), 2002)

"The tyranny of scales will not be simply defeated by building bigger and faster computers" (SBES report, p. 30)

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Requirements and impact of algorithms?

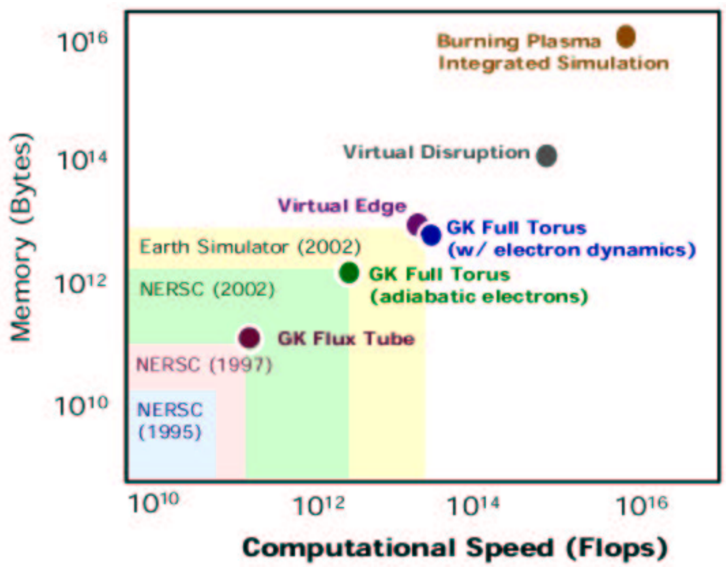

(c) Computational requirements for fusion plasma integrated simulation (FSP report)

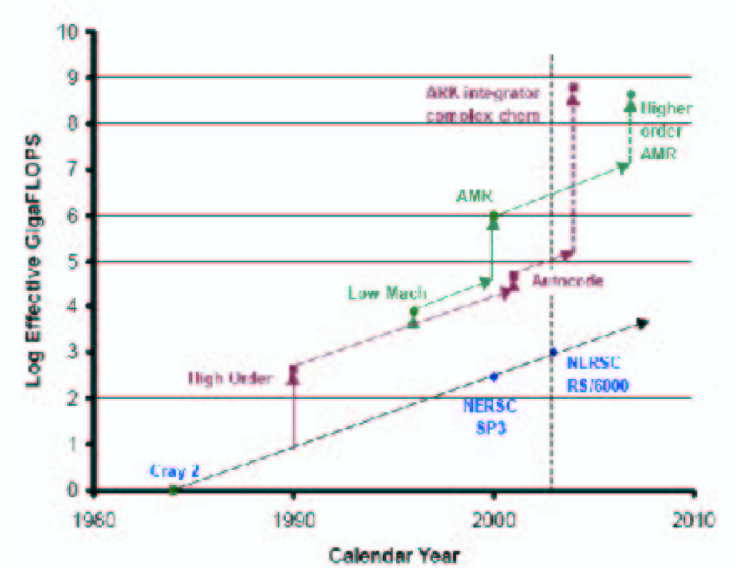

(d) Impact of algorithms in gas combustion effective FLOPS (SBES report)

"Faster and more cost-effective hardware is a strong driver for simulation-based engineering. However, algorithmic improvements have been far more important." (SBES report, p. 50)

Luis Chacón, chacon@lanl.gov

Algorithmic challenges in XMHD (I)

- XMHD is a strongly hyperbolic PDE system.
- Numerically, XMHD is a nonlinear algebraic system of very stiff equations:
  - Elliptic stiffness (diffusion): Jacobian condition number $\sim \frac{\Delta t\, D}{\Delta x^2}$
  - Hyperbolic stiffness (linear and dispersive waves): Jacobian condition number $\sim \Delta t\, \omega_{fast} \sim \frac{\Delta t}{\Delta t_{CFL}} \gg 1$
- An implicit integration of XMHD *may be* advantageous to step over wave phenomena and get to the dynamical time scale of interest.

Implicit methods require inversion of very large, sparse matrices!

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Algorithmic challenges in XMHD (II)

- Brute-force algorithms will not be able to cover the span between disparate time/length scales, regardless of computer power (gas combustion example).
- Key algorithmic feature: SCALABILITY!
  - Minimize number of degrees of freedom (grid points) without sacrificing spatial resolution: spatial adaptivity
  - Be able to follow lowest frequency time scales (application dependent): implicit time stepping

It is our contention that fully implicit, spatially adaptive methods are essential for scalability, and thus an integral part of a predictive plasma simulation tool!

Scalable algorithm: CPU $\sim \mathcal{O}(N/np)$, N is # dof, np is # procs

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Alternatives for inversion algorithms: NOT SCALABLE

- Explicit (trivial option, here for comparison):

$$CPU \sim \frac{T_{max}}{\Delta t_{CFL}} \times \frac{N}{np} \sim \mathcal{O}(N^{1+\alpha/d}/np) \ ;$$

d is # dimensions, $\alpha = 1, 2$ for linear, dispersive waves

- Implicit (requires matrix inversion). Naive options are NOT scalable:
  - Direct methods: good parallelization, but do not scale with problem size:

$$CPU \sim \mathcal{O}\left(\frac{N^{(3-2/d)}}{np^{\beta}}\right) \ , \ \beta \gtrsim 1$$

  - Iterative methods (unpreconditioned Krylov, stationary, etc.): good paralellization, but VERY slow convergence:

$$CPU \sim \mathcal{O}\left(\frac{N^{\alpha}}{np^{\beta}}\right) \ , \ \alpha > 1 \ ; \ \beta \gtrsim 1$$

Scalingwise, direct solver is WORSE than explicit for $d > 1 + \frac{\alpha}{2} \approx 2$

Alternatives for inversion algorithms: SCALABLE

- Scalable matrix inversion methods require MULTILEVEL approaches (divide and conquer in wavenumbers):
  - Direct-solve substructuring (X. Z. Tang).
  - FETI-DP (Glasser's talk).
  - Multilevel iterative (e.g., classical MG, algebraic MG).

$$CPU \sim \mathcal{O}\left(\frac{N\ \log(N)}{np^{\beta}}\right)\ ,\ \beta \gtrsim 1$$

- Both approaches are being pursued in T-15 at LANL!
- This talk focuses on the second approach.
- A fundamental component of iterative ML methods (both classical and algebraic) is the existence of a SMOOTHER (convergent stationary iterative method).

Q: How to ensure the existence of a SMOOTHER for XMHD? A: Parabolization!

Luis Chacón, chacon@lanl.gov

Parabolization and Schur complement: an example

- PARABOLIZATION EXAMPLE:

$$\partial_t u = \partial_x v \ , \ \partial_t v = \partial_x u.$$

$$\begin{aligned} u^{n+1} &= u^n + \Delta t \partial_x v^{n+1}, \\ v^{n+1} &= v^n + \Delta t \partial_x u^{n+1}. \end{aligned}$$

$$\boxed{(I - \Delta t^2 \partial_{xx}) u^{n+1} = u^n + \Delta t \partial_x v^n}$$

- PARABOLIZATION via SCHUR COMPLEMENT:

$$\begin{bmatrix} D_1 & U \\ L & D_2 \end{bmatrix} = \begin{bmatrix} I & U D_2^{-1} \\ 0 & I \end{bmatrix} \begin{bmatrix} D_1 - U D_2^{-1} L & 0 \\ 0 & D_2 \end{bmatrix} \begin{bmatrix} I & 0 \\ D_2^{-1} L & I \end{bmatrix}.$$

Stiff off-diagonal blocks L, U now sit in diagonal via Schur complement $D_1 - U D_2^{-1} L$. The system has been "PARABOLIZED."

$$D_1 - U D_2^{-1} L = (I - \Delta t^2 \partial_{xx})$$

How to build a successful fully implicit algorithm for XMHD

- Even if a smoother exists, MG is remarkably temperamental
- Combination of Krylov methods and MG is optimal:
  - MG provides scalability (as a preconditioner)
  - Krylov provides robustness
- We seek to develop a successful algorithm for XMHD based on Newton-Krylov-MG.
- We will start with resistive MHD, and then move to XMHD.
- Finally we will discuss the combination of implicit time stepping with dynamic grid adaptation.

Luis Chacón, chacon@lanl.gov

Jacobian-Free Newton-Krylov Methods

- Objective: solve nonlinear system $\vec{G}(\vec{x}^{n+1}) = \vec{0}$ efficiently.
- Converge nonlinear couplings using Newton-Raphson method: $\left.\frac{\partial \vec{G}}{\partial \vec{x}}\right|_k \delta \vec{x}_k = -\vec{G}(\vec{x}_k)$.
- Jacobian-free implementation: $\left(\frac{\partial \vec{G}}{\partial \vec{x}}\right)_k \vec{y} = J_k \vec{y} = \lim_{\epsilon \to 0} \frac{\vec{G}(\vec{x}_k + \epsilon \vec{y}) - \vec{G}(\vec{x}_k)}{\epsilon}$
- Krylov method of choice: GMRES (nonsymmetric systems).
- Right preconditioning: solve equivalent Jacobian system for $\delta y = P_k \delta \vec{x}$:

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

APPROXIMATIONS IN PRECONDITIONER DO NOT AFFECT ACCURACY OF CONVERGED SOLUTION; THEY ONLY AFFECT EFFICIENCY!

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Implicit *resistive* MHD solver

Resistive MHD model equations

$$
\begin{aligned}
\frac{\partial \rho}{\partial t} &+ \nabla \cdot (\rho \vec{v}) = 0, \\
\frac{\partial \vec{B}}{\partial t} &+ \nabla \times \vec{E} = 0, \\
\frac{\partial (\rho \vec{v})}{\partial t} + \nabla \cdot \left[\rho \vec{v} \vec{v} - \vec{B} \vec{B} \right. &- \left. \rho \nu \nabla \vec{v} + \overleftrightarrow{I} (p + \frac{B^2}{2}) \right] = 0, \\
\frac{\partial T}{\partial t} + \vec{v} \cdot \nabla T &+ (\gamma - 1) T \nabla \cdot \vec{v} = 0,
\end{aligned}
$$

- Plasma is assumed polytropic $p \propto n^{\gamma}$.
- Resistive Ohm's law:

$$
\vec{E} = -\vec{v} \times \vec{B} + \eta \nabla \times \vec{B}
$$

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Resistive MHD Jacobian block structure

- The linearized resistive MHD model has the following couplings:

$$\begin{aligned}
\delta\rho &= L_\rho(\delta\rho, \delta\vec{v}) \\
\delta T &= L_T(\delta T, \delta\vec{v}) \\
\delta\vec{B} &= L_B(\delta\vec{B}, \delta\vec{v}) \\
\delta\vec{v} &= L_v(\delta\vec{v}, \delta\vec{B}, \delta\rho, \delta T)
\end{aligned}$$

- Therefore, the Jacobian of the resistive MHD model has the following coupling structure:

$$J\delta\vec{x} = \begin{bmatrix} D_\rho & 0 & 0 & U_{v\rho} \\ 0 & D_T & 0 & U_{vT} \\ 0 & 0 & D_B & U_{vB} \\ L_{\rho v} & L_{Tv} & L_{Bv} & D_v \end{bmatrix} \begin{pmatrix} \delta\rho \\ \delta T \\ \delta\vec{B} \\ \delta\vec{v} \end{pmatrix}$$

- Diagonal blocks contain advection-diffusion contributions, and are “easy” to invert using MG techniques. Off diagonal blocks L and U contain all hyperbolic couplings.

Luis Chacón, chacon@lanl.gov

PARABOLIZATION: Schur complement formulation

- We consider the block structure:

$$J\delta\vec{x} = \begin{bmatrix} M & U \\ L & D_v \end{bmatrix} \begin{pmatrix} \delta\vec{y} \\ \delta\vec{v} \end{pmatrix}$$

$$\delta\vec{y} = \begin{pmatrix} \delta\rho \\ \delta T \\ \delta\vec{B} \end{pmatrix} \quad ; \quad M = \begin{pmatrix} D_\rho & 0 & 0 \\ 0 & D_T & 0 \\ 0 & 0 & D_B \end{pmatrix}$$

- M is "easy" to invert (advection-diffusion, MG-friendly).

Schur complement analysis of 2x2 block J yields:

$$\begin{bmatrix} M & U \\ L & D_v \end{bmatrix}^{-1} = \begin{bmatrix} I & 0 \\ -LM^{-1} & I \end{bmatrix} \begin{bmatrix} M^{-1} & 0 \\ 0 & P_{Schur}^{-1} \end{bmatrix} \begin{bmatrix} I & -M^{-1}U \\ 0 & I \end{bmatrix},$$

$$\boxed{P_{Schur} = D_v - LM^{-1}U}.$$

- EXACT Jacobian inverse only requires M^{-1} and P_{Schur}^{-1}.
- Schur complement formulation is fundamentally unchanged in Hall MHD!

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Physics-based preconditioner (I)

- The Schur complement analysis translates into the following 3-step EXACT inversion algorithm:

$$\begin{aligned}
\text{Predictor} &: \quad \delta\vec{y}^* = -M^{-1}G_y \\
\text{Velocity update} &: \quad \delta\vec{v} = P_{Schur}^{-1}[-G_v - L\delta\vec{y}^*], \;\; P_{Schur} = D_v - LM^{-1}U \\
\text{Corrector} &: \quad \delta\vec{y} = \delta\vec{y}^* - M^{-1}U\delta\vec{v}
\end{aligned}$$

- MG treatment of P_{Schur} is impractical due to M^{-1}.

$$\boxed{\text{Need suitable simplifications (SEMI-IMPLICIT)!}}$$

- We consider the small-flow-limit case: $M^{-1} \approx \Delta t$
- This approximation is equivalent to splitting flow in original equations.

Los Alamos
NATIONAL LABORATORY
EST. 1943

Physics-based preconditioner (II)

- Small flow approximation: $M^{-1} \approx \Delta t$ in steps 2 & 3 of Schur algorithm:

$$\begin{aligned} \delta \vec{y}^* &= -M^{-1} G_y \\ \delta \vec{v} &\approx P_{SI}^{-1} [-G_v - L \delta \vec{y}^*] \; ; \; P_{SI} = D_v - \Delta t L U \\ \delta \vec{y} &\approx \delta \vec{y}^* - \Delta t U \delta \vec{v} \end{aligned}$$

where:

$$\boxed{P_{SI} = \rho^n \left[\overleftrightarrow{I} / \Delta t + \theta (\vec{v}_0 \cdot \nabla \overleftrightarrow{I} + \overleftrightarrow{I} \cdot \nabla \vec{v}_0 - \nu^n \nabla^2 \overleftrightarrow{I}) \right] + \Delta t \theta^2 W(\vec{B}_0, p_0)}$$

$$W(\vec{B}_0, p_0) = \vec{B}_0 \times \nabla \times \nabla \times [\overleftrightarrow{I} \times \vec{B}_0] - \vec{j}_0 \times \nabla \times [\overleftrightarrow{I} \times \vec{B}_0] - \nabla [\overleftrightarrow{I} \cdot \nabla p_0 + \gamma p_0 \nabla \cdot \overleftrightarrow{I}]$$

- P_{SI} is block diagonally dominant by construction!
- We employ multigrid methods (MG) to approximately invert P_{SI} and M: 1 V(4,4) cycle

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Efficiency: Δt scaling (2D tearing mode)

32×32

Δt	Newton/Δt	GMRES/Δt	CPU (s)	CPU_{exp}/CPU	$\Delta t/\Delta t_{CFL}$
2	5.9	20.9	115	3.1	354
3	5.9	25.6	139	3.8	531
4	6.0	30.5	163	4.3	708
6	6.0	34.7	184	5.8	1062

128×128

Δt	Newton/Δt	GMRES/Δt	CPU (s)	CPU_{exp}/CPU	$\Delta t/\Delta t_{CFL}$
0.5	4.9	8.4	764	8.0	380
0.75	5.7	10.2	908	10.0	570
1.0	5.0	11.5	1000	12.7	760
1.5	5.6	14.7	1246	14.6	1140

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Efficiency: grid scaling

$\Delta t \approx 1100 \Delta t_{CFL}$, 10 time steps

Grid	Δt	Newton/Δt	GMRES/Δt	CPU	$\widehat{CPU}$
32x32	6	6.0	34.7	184	5.3
64x64	3	5.8	22.9	468	20.4
128x128	1.5	5.6	14.8	1246	84.2

Why does GMRES/Δt decrease with resolution?

Los Alamos NATIONAL LABORATORY EST. 1943

Effect of spatial truncation error

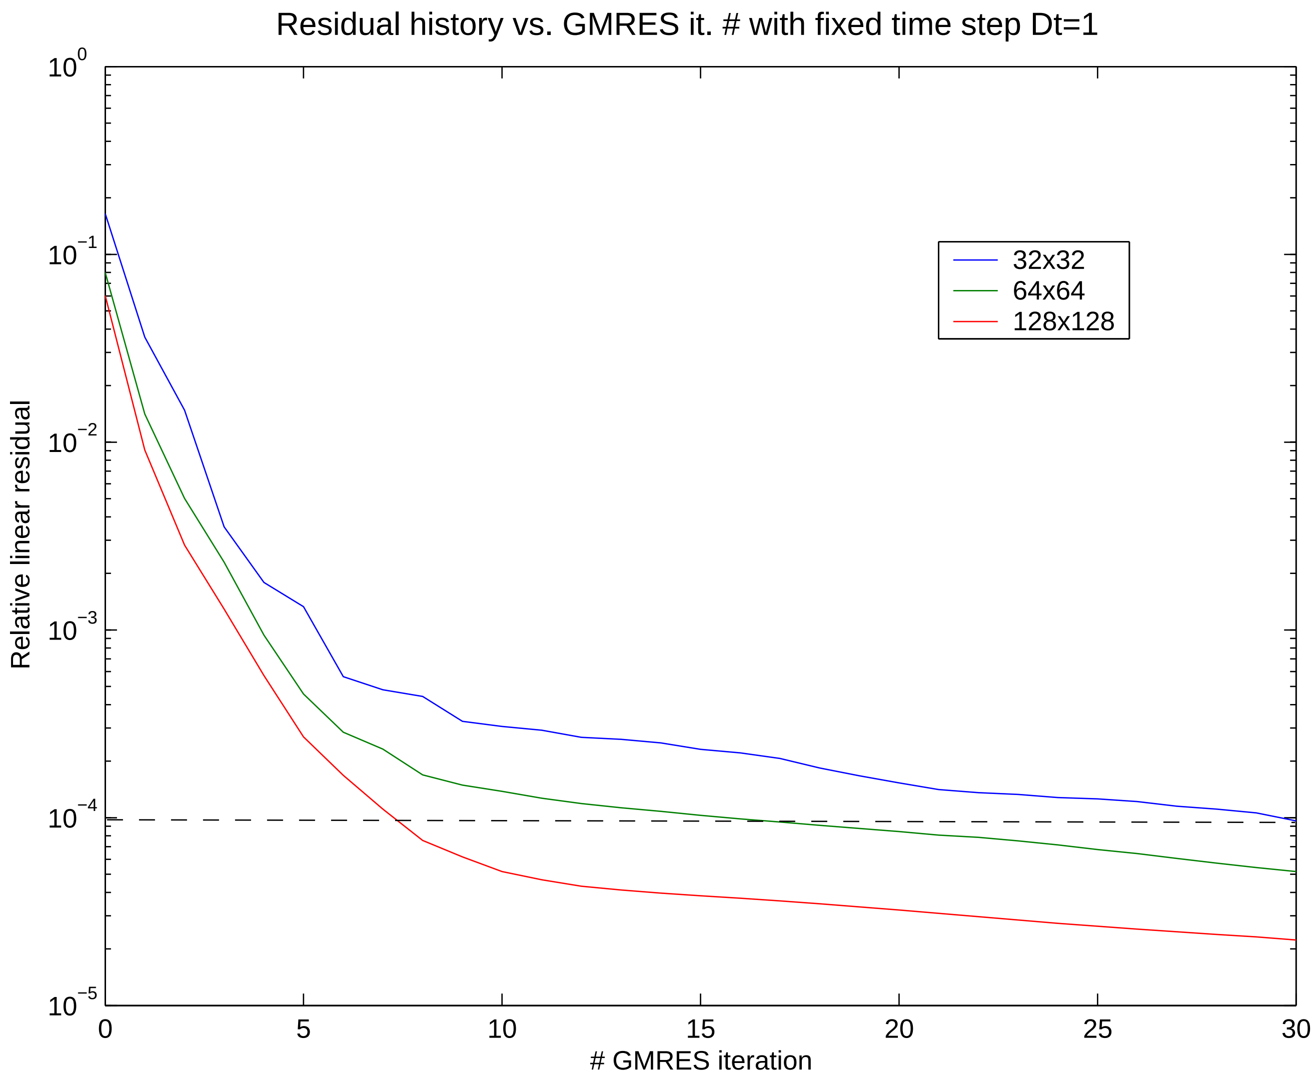

Luis Chacón, chacon@lanl.gov

Sample 3D results: Screw pinch in 3D

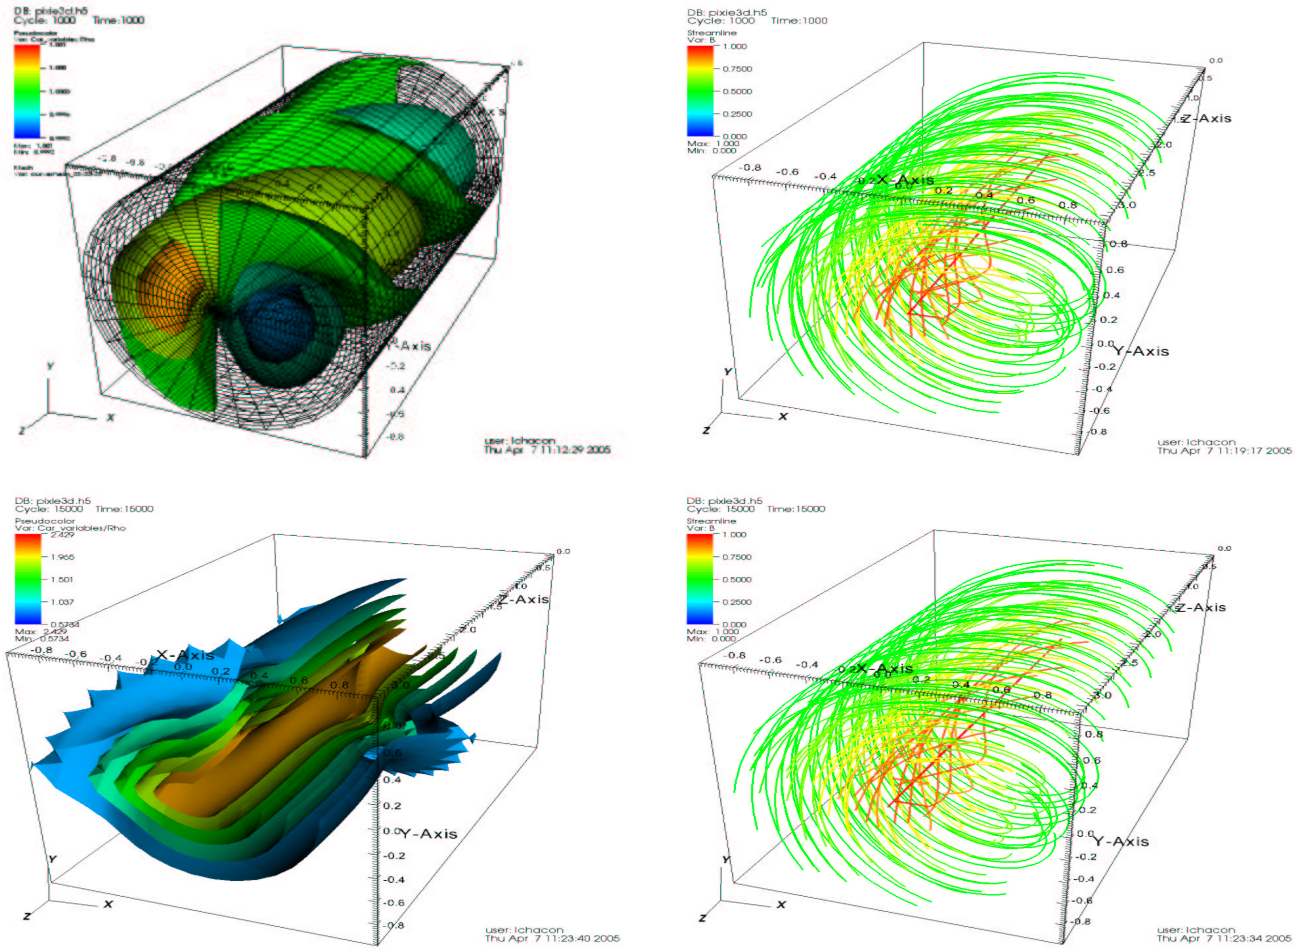

Luis Chacón, chacon@lanl.gov

Sample 3D results: 3D KHI

Knoll and Brackbill, Phys. Plasmas 9 (9) 2002

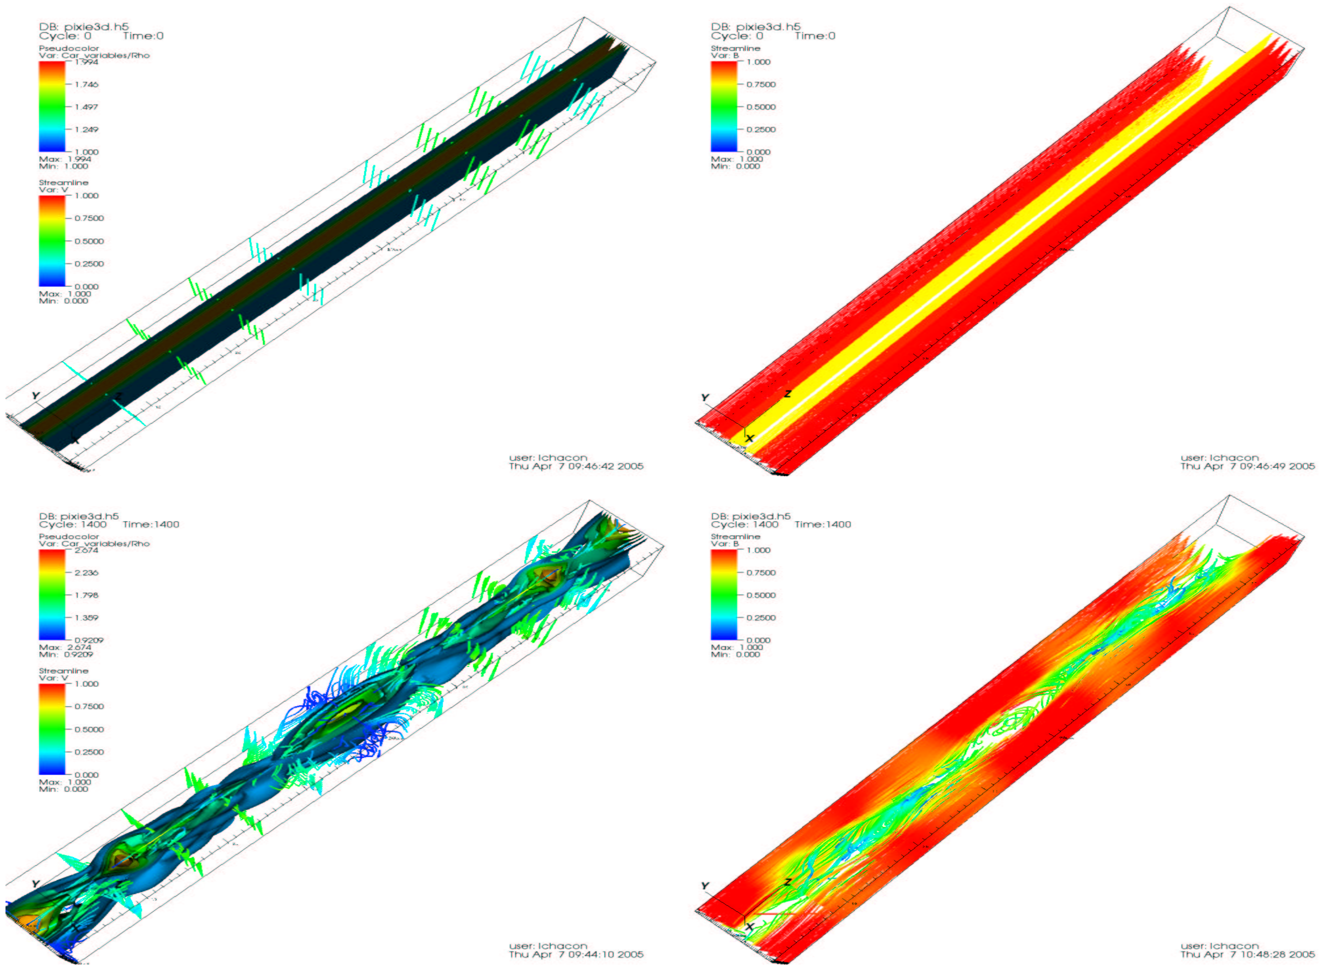

Implicit *extended* MHD solver

Extended MHD model equations

$$\begin{aligned}
\frac{\partial \rho}{\partial t} &+ \nabla \cdot (\rho \vec{v}) = 0, \\
\frac{\partial \vec{B}}{\partial t} &+ \nabla \times \vec{E} = 0, \\
\frac{\partial (\rho \vec{v})}{\partial t} + \nabla \cdot \left[\rho \vec{v}\vec{v} - \vec{B}\vec{B} \right. &- \left. \rho \nu \nabla \vec{v} + \overleftrightarrow{I}(p + \frac{B^2}{2}) \right] = 0, \\
\frac{\partial T_e}{\partial t} + \vec{v} \cdot \nabla T_e &+ (\gamma - 1) T_e \nabla \cdot \vec{v} = 0,
\end{aligned}$$

- Plasma is assumed polytropic $p \propto n^\gamma$.
- We assume cold ion limit: $T_i \ll T_e \Rightarrow \boxed{p \approx p_e}$.
- Generalized Ohm's law:

$$\boxed{\vec{E} = -\vec{v} \times \vec{B} + \eta \nabla \times \vec{B} - \frac{d_i}{\rho}(\vec{j} \times \vec{B} - \nabla p_e)}$$

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Extended MHD Jacobian block structure

- The linearized extended MHD model has the following couplings:

$$
\begin{aligned}
\delta\rho &= L_\rho(\delta\rho, \delta\vec{v}) \\
\delta T &= L_T(\delta T, \delta\vec{v}) \\
\delta\vec{B} &= L_B(\delta\vec{B}, \delta\vec{v}, \delta\rho, \delta T) \\
\delta\vec{v} &= L_v(\delta\vec{v}, \delta\vec{B}, \delta\rho, \delta T)
\end{aligned}
$$

- Jacobian coupling structure:

$$
J\delta\vec{x} = \begin{bmatrix} D_\rho & 0 & 0 & U_{v\rho} \\ 0 & D_T & 0 & U_{vT} \\ L_{\rho B} & L_{TB} & D_B & U_{vB} \\ L_{\rho v} & L_{Tv} & L_{Bv} & D_v \end{bmatrix} \begin{pmatrix} \delta\rho \\ \delta T \\ \delta\vec{B} \\ \delta\vec{v} \end{pmatrix}
$$

- We have added off-diagonal couplings.

Luis Chacón, chacon@lanl.gov

Extended MHD Jacobian block structure (cont.)

- The coupling structure can be substantially simplified if we note $(p \approx p_e)$:

$$\frac{1}{\rho}(\vec{j} \times \vec{B} - \nabla p_e) \approx \frac{D\vec{v}}{Dt}$$

and therefore:

$$\vec{E} \approx -\vec{v} \times \vec{B} + \frac{\eta(T)}{\mu_0} \nabla \times \vec{B} - d_i \frac{D\vec{v}}{Dt}$$

- This transforms jacobian coupling structure to:

$$J\delta\vec{x} \approx \begin{bmatrix} D_\rho & 0 & 0 & U_{v\rho} \\ 0 & D_T & 0 & U_{vT} \\ 0 & 0 & D_B & U_{vB}^R + U_{vB}^H \\ L_{\rho v} & L_{Tv} & L_{Bv} & D_v \end{bmatrix} \begin{pmatrix} \delta\rho \\ \delta T \\ \delta\vec{B} \\ \delta\vec{v} \end{pmatrix}$$

We can therefore reuse ALL resistive MHD PC framework!

Luis Chacón, chacon@lanl.gov

Extended MHD preconditioner

- Use same Schur complement approach.
- M block contains ion scales only! Approximation $M^{-1} \approx \Delta t$ is very good in extended MHD (ion scales do NOT contribute to numerical stiffness).
- Additional block U_{vB}^{H} results, after the Schur complement treatment, in systems of the form:

$$\partial_t \delta\vec{v} - d_i \vec{B}_0 \times (\nabla \times \nabla \times \delta\vec{v}) = rhs$$

- This system supports dispersive waves $\omega \sim k^2$!
- We have shown analytically that damped JB is a smoother for these systems!

We can use classical MG!

Luis Chacón, chacon@lanl.gov

Preliminary efficiency results (2D tearing mode)

$d_i = 0.05$

1 time step, $\Delta t = 1.0$, V(3,3) cycles, mg_tol=1e-2

Grid	Newton/Δt	GMRES/Δt	CPU (s)	CPU_{exp}/CPU	$\Delta t/\Delta t_{exp}$
32x32	5	22	25	0.44	110
64x64	5	12	66	1.4	238
128x128	5	8	164	6.2	640
256x256	4	7	674	30	3012

Again, GMRES/Δt decreases with resolution!

Luis Chacón, chacon@lanl.gov

Effect of spatial truncation error

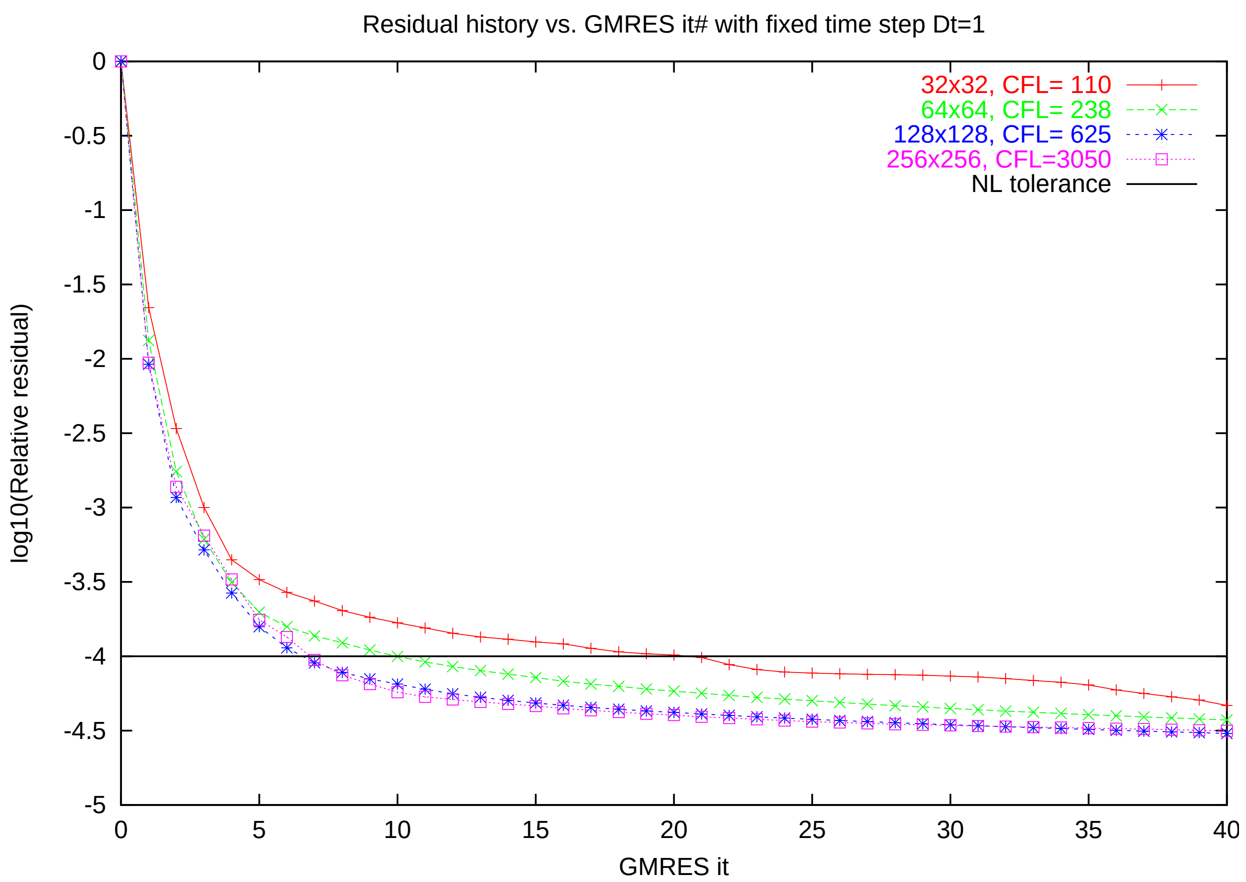

Luis Chacón, chacon@lanl.gov

GEM Challenge

J. Birn et al., *J. Geophys. Res.*, 106 (A3), p.3715-19 (2001)

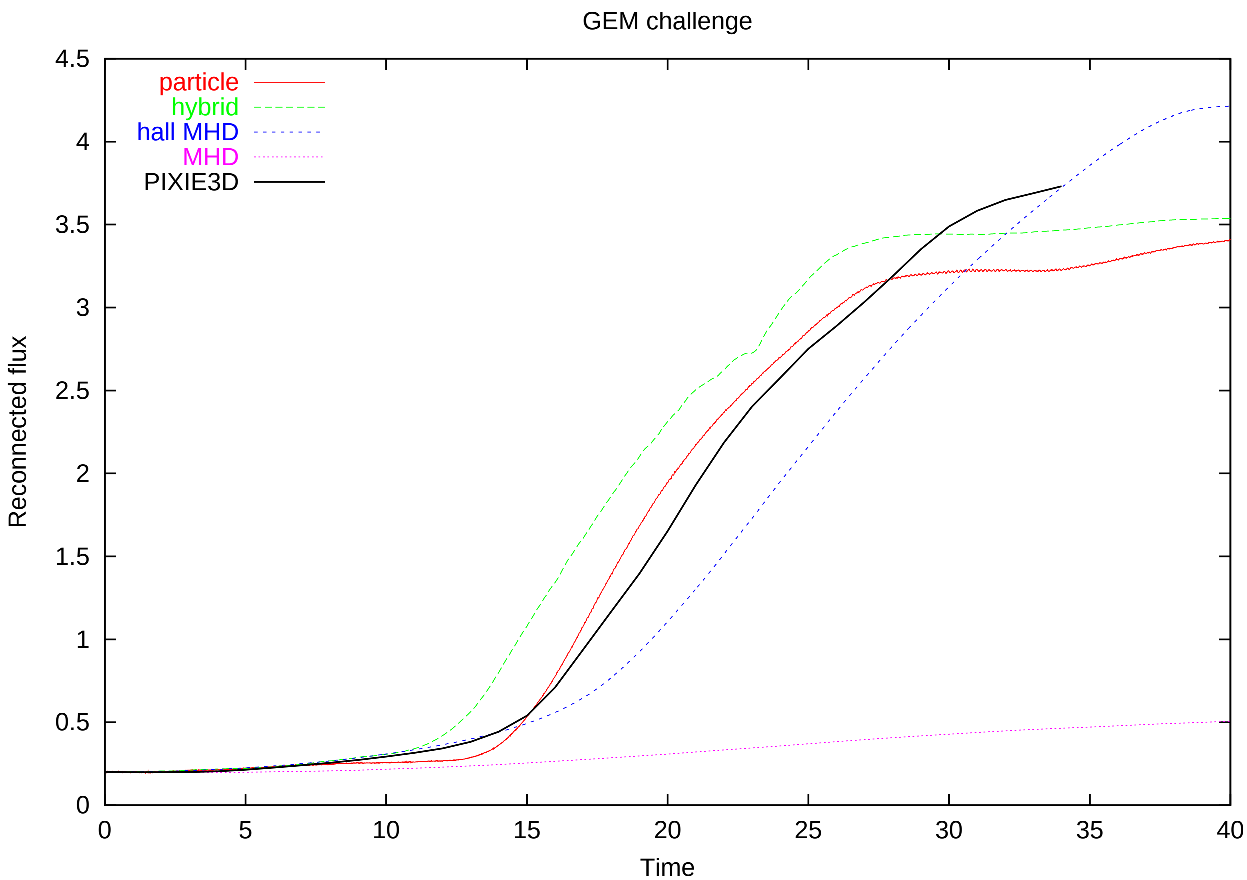

Luis Chacón, chacon@lanl.gov

Parallel performance with PETSc Toolkit (unpreconditioned)

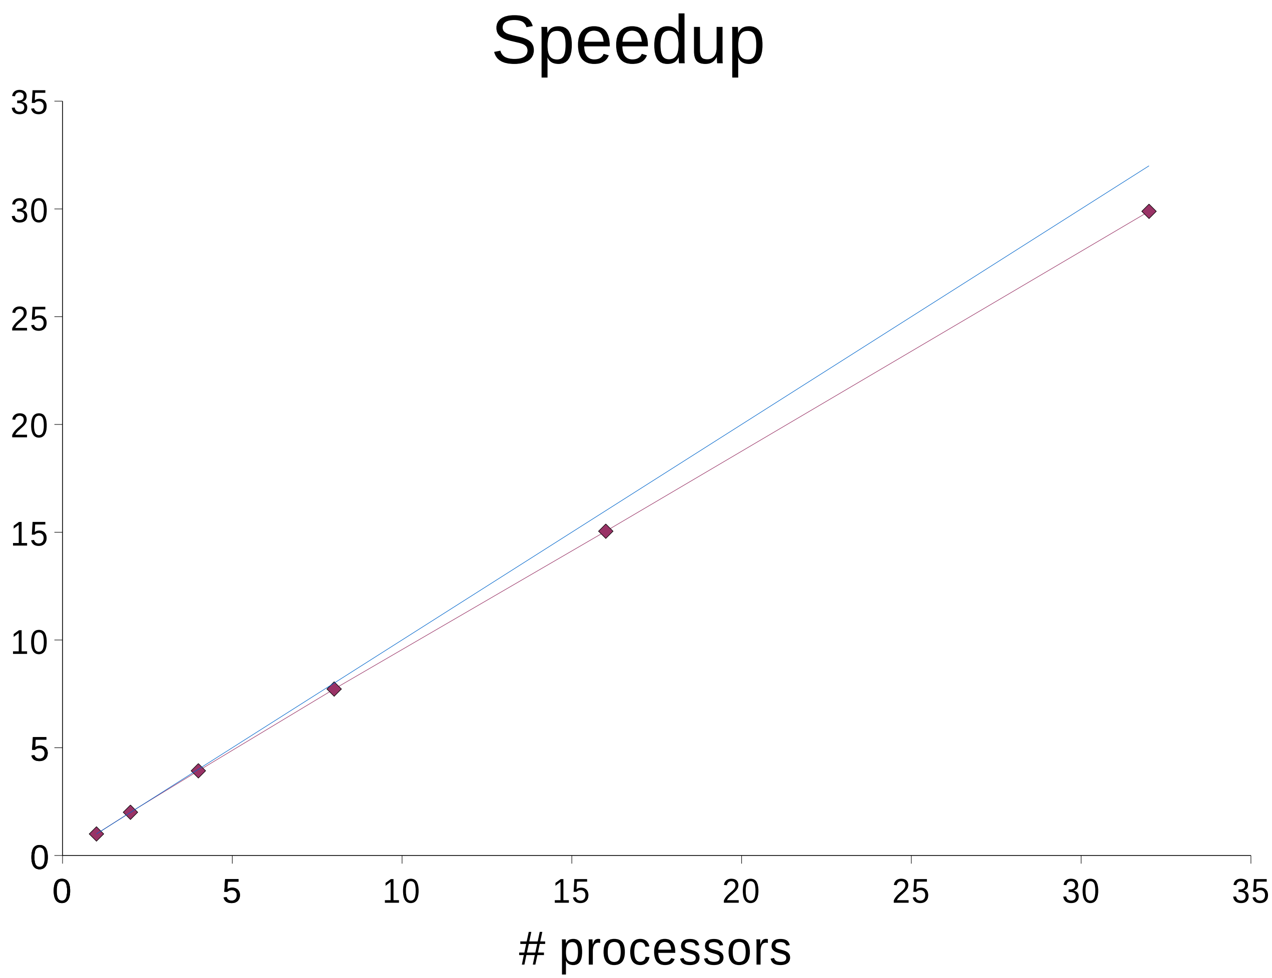

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Implicit-AMR proof of principle

B. Philip, M. Pernice, and L. Chacón, *Lecture Notes in Computational Science and Engineering*, accepted (2006).

Current-Vorticity Formulation of Reduced Resistive MHD[1]

$$
\begin{aligned}
(\partial_t + \mathbf{u}\cdot\nabla - \eta\Delta)\, J + \Delta E_0 &= \mathbf{B}\cdot\nabla\omega + \{\Phi,\Psi\} \\
(\partial_t + \mathbf{u}\cdot\nabla - \nu\Delta)\, \omega + S_\omega &= \mathbf{B}\cdot\nabla J \\
\Delta\Phi &= \omega \\
\Delta\Psi &= J
\end{aligned}
$$

$$
\mathbf{u} = \vec{z}\times\nabla\Phi \ , \ \mathbf{B} = \vec{z}\times\nabla\Psi
$$

$$
\{\Phi,\Psi\} = 2[\Phi_{xy}(\Psi_{xx}-\Psi_{yy}) - \Psi_{xy}(\Phi_{xx}-\Phi_{yy})]
$$

Preconditioner is developed as an extension of
Chacón, Knoll and Finn, JCP, **178** (2002).

[1] Strauss and Longcope, JCP, **147**, 1998

Structured Adaptive Mesh Refinement

- *Structured* adaptive mesh refinement (SAMR) represents a locally refined mesh as a union of logically rectangular meshes.

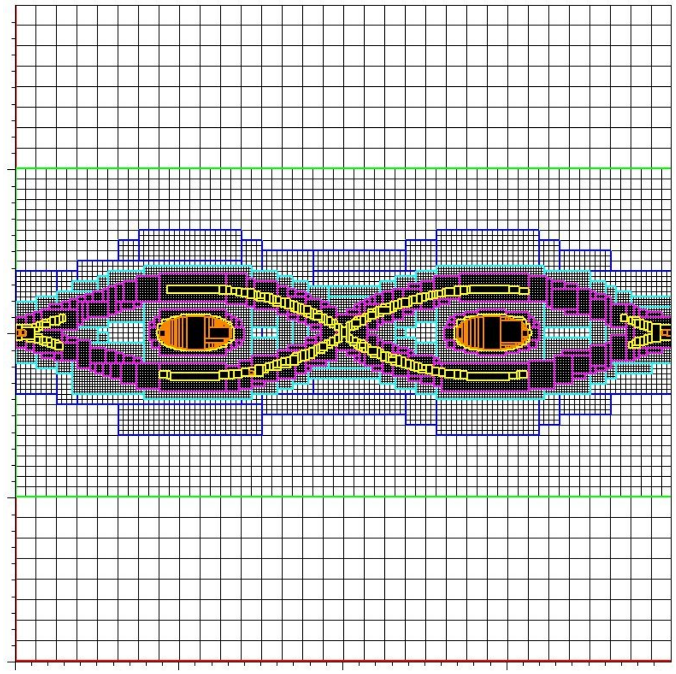

- The mesh is organized as a hierarchy of nested *refinement levels*.
- Each refinement level defines a region of uniform resolution.
- Each refinement level is the union of logically rectangular *patches*.

Luis Chacón, chacon@lanl.gov

Hierarchical Structure of SAMR Grids

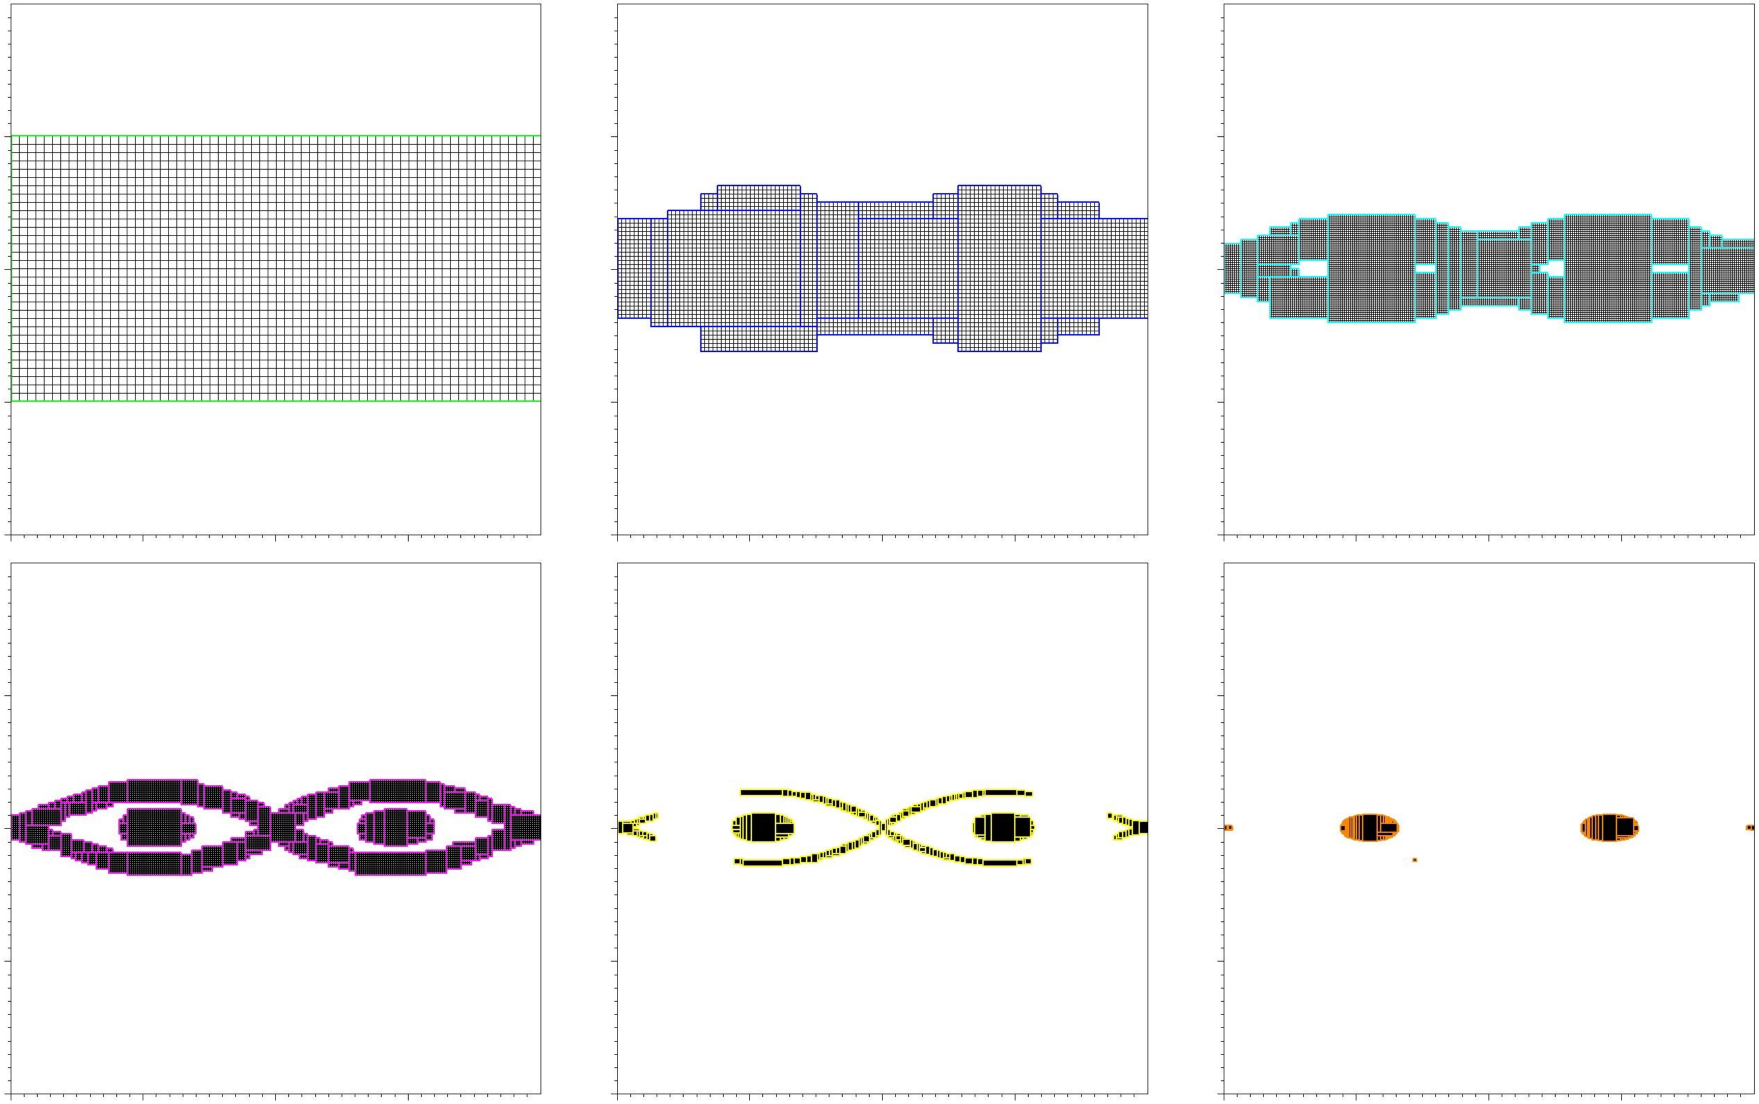

Luis Chacón, chacon@lanl.gov

Tearing Mode Results

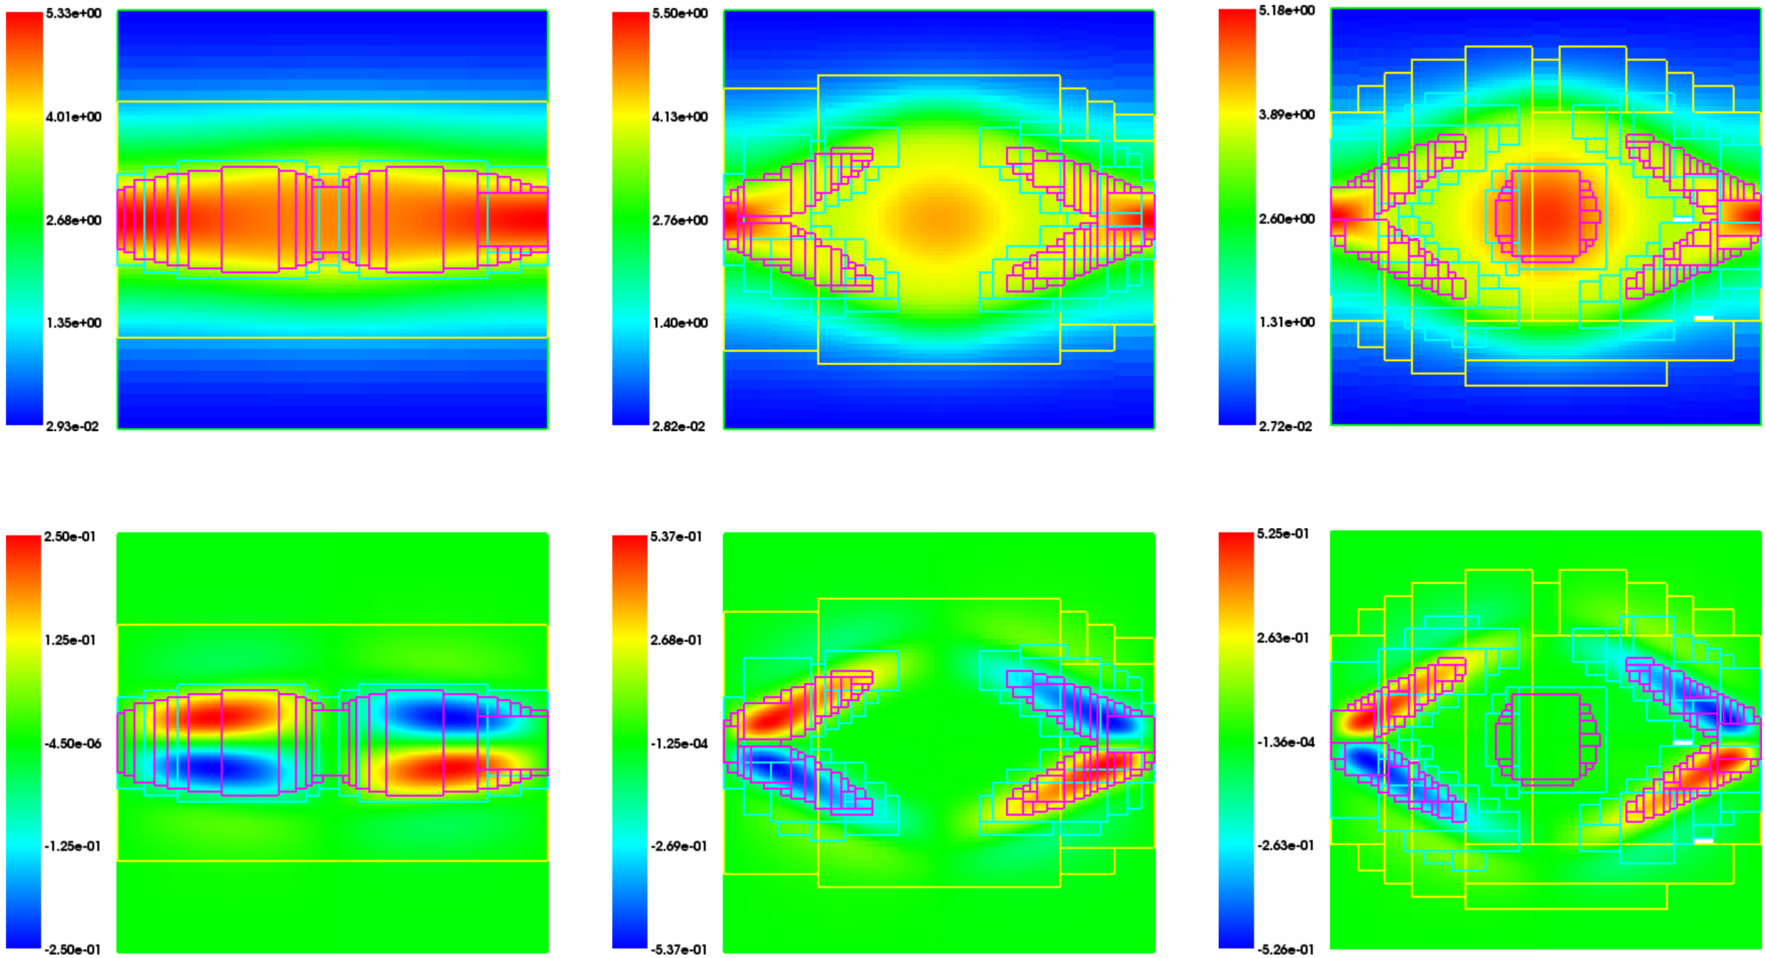

$t = 50$ $\qquad$ $t = 120$ $\qquad$ $t = 200$

Los Alamos
NATIONAL LABORATORY
EST. 1943

Luis Chacón, chacon@lanl.gov

Tearing Mode Performance

	NNI					NLI				
Levels	1	2	3	4	5	1	2	3	4	5
32×32	1.5	2.0	2.0	2.1	2.5	3.4	7.9	12.0	19.3	33.7
64×64	1.8	2.0	2.0	2.4	–	6.5	11.7	19.1	33.2	–
128×128	1.8	2.0	2.4	–	–	12.5	20.1	27.2	–	–
256×256	1.9	2.0	–	–	–	19.9	27.5	–	–	–
512×512	1.9	–	–	–	–	26.3	–	–	–	–

$\eta_k = 0.1$, $\epsilon_{rel} = \epsilon_{abs} = 10^{-7}$, 2 SI iterations, V(3,3) cycles

Los Alamos NATIONAL LABORATORY EST. 1943

Luis Chacón, chacon@lanl.gov

Island Coalescence Results at t=8

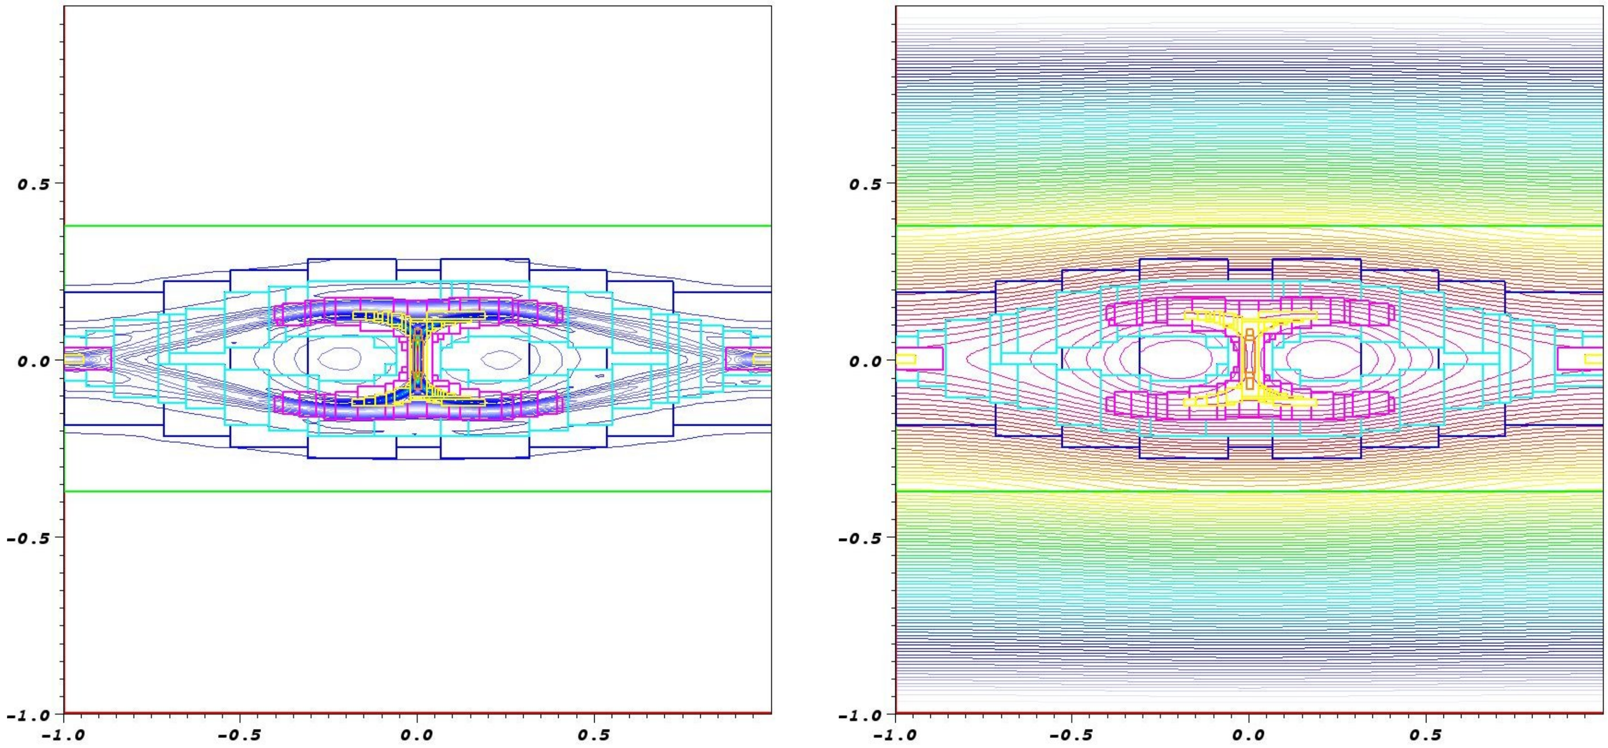

Luis Chacón, chacon@lanl.gov

Tilt Instability Results at t=7