Progress on Scalable Parallel Computation for Extended MHD Modeling of Fusion Plasmas

Alan H. Glasser
Collaborator: V. S. Lukin

Presented at the 50th Annual Meeting of the
American Physical Society, Division of Plasma Physics
Dallas, Texas November 17-21, 2008

Scalability of Extended MHD Simulation

- 3D extended MHD modeling of magnetically confined fusion plasmas requires petascale computing: 1 petaflop = 10^{15} flops ~10^5 procs.
- Scalability: doubling problem size and number of processors causes little or no change in cpu time to solution.
- Advanced extended MHD codes use high-order methods of spatial discretization. NIMROD, M3D, SEL/HiFi.
- Known scalable methods for elliptic and parabolic systems:
  - Multigrid. Applicable to low-order spatial discretization.
  - FETI-DP domain substructuring. Applicable to high-order spatial discretization.
- Extended MHD dominated by hyperbolic waves, multiple time scales. Requires parabolization, physics-based preconditioning. Luis Chacon.
- Matrix-free Newton-Krylov iteration.

Organization of Presentation

- The SEL/HiFi spectral element code
- Physics-based preconditioning.
- Preconditioners for ideal and Hall MHD.
- Static condensation and FETI-DP.
- Scaling results.
- Future plans.

SEL/HiFi Spectral Element Code

- Flux-source form: simple, general problem setup.
- Spatial discretization:
  - High-order C^0 spectral elements, modal basis
  - Harmonic grid generation, adaptation, alignment
- Time step: fully implicit, 2nd-order accurate,
  - θ-scheme
  - BDF2
- Static condensation, Schur complement.
  - Small local direct solves for grid cell interiors.
  - Preconditioned GMRES for Schur complement.
- Distributed parallel operation with MPI and PETSc.

Spatial Discretization

Flux-Source Form of Equations

$$\frac{\partial u^i}{\partial t} + \nabla \cdot \mathbf{F}^i = S^i$$

$$\mathbf{F}^i = \mathbf{F}^i(t, \mathbf{x}, u^j, \nabla u^j)$$

$$S^i = S^i(t, \mathbf{x}, u^j, \nabla u^j)$$

Galerkin Expansion

$$u^i(t, \mathbf{x}) \approx \sum_{j=0}^{n} u^i_j(t) \alpha_j(\mathbf{x})$$

Weak Form of Equations

$$(\alpha_i, \alpha_j) \dot{u}^k_j = \int_\Omega d\mathbf{x} \left(S^k \alpha_i + \mathbf{F}^k \cdot \nabla \alpha_i \right) - \int_{\partial\Omega} d\mathbf{x} \alpha_i \mathbf{F}^k \cdot \hat{\mathbf{n}}$$

Physics-Based Preconditioning

Factorization and Schur Complement

Linear System

$$\mathbf{L}\mathbf{u} = \mathbf{r}, \quad \mathbf{L} \equiv \begin{pmatrix} \mathbf{L}_{11} & \mathbf{L}_{12} \\ \mathbf{L}_{21} & \mathbf{L}_{22} \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \end{pmatrix}, \quad \mathbf{r} = \begin{pmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{pmatrix}$$

Factorization

$$\mathbf{L} \equiv \begin{pmatrix} \mathbf{L}_{11} & \mathbf{L}_{12} \\ \mathbf{L}_{21} & \mathbf{L}_{22} \end{pmatrix} = \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{L}_{21}\mathbf{L}_{11}^{-1} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{L}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{S} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{L}_{11}^{-1}\mathbf{L}_{12} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}$$

Schur Complement

$$\mathbf{S} \equiv \mathbf{L}_{22} - \mathbf{L}_{21}\mathbf{L}_{11}^{-1}\mathbf{L}_{12}$$

Exact and Approximate Inverse

Preconditioned Krylov Iteration

Inverse

$$\mathsf{L}^{-1} = \begin{pmatrix} \mathsf{I} & -\mathsf{L}_{11}^{-1}\mathsf{L}_{12} \\ \mathsf{0} & \mathsf{I} \end{pmatrix} \begin{pmatrix} \mathsf{L}_{11}^{-1} & \mathsf{0} \\ \mathsf{0} & \mathsf{S}^{-1} \end{pmatrix} \begin{pmatrix} \mathsf{I} & \mathsf{0} \\ -\mathsf{L}_{21}\mathsf{L}_{11}^{-1} & \mathsf{I} \end{pmatrix}$$

Exact Solution

$$\mathbf{s}_1 = \mathsf{L}_{11}^{-1}\mathbf{r}_1, \quad \mathbf{s}_2 = \mathbf{r}_2 - \mathsf{L}_{21}\mathbf{s}_1$$

$$\mathbf{u}_2 = \mathsf{S}^{-1}\mathbf{s}_2, \quad \mathbf{u}_1 = \mathbf{s}_1 - \mathsf{L}_{11}^{-1}\mathsf{L}_{12}\mathbf{u}_2$$

Preconditioned Krylov Iteration

$$\mathsf{P} \approx \mathsf{L}^{-1}, \quad (\mathsf{L}\mathsf{P})\left(\mathsf{P}^{-1}\mathbf{u}\right) = \mathbf{r}$$

Outer iteration preserves full nonlinear accuracy.
Need approximate Schur complement **S**
and scalable solution procedure for $\mathbf{L}_{11}$ and **S**.

Ideal MHD Waves

Linearized, Normalized Equations

$$\frac{\partial p}{\partial t} + \gamma \nabla \cdot \mathbf{v} = 0, \quad \frac{\partial \mathbf{b}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B})$$

$$\frac{\partial \mathbf{v}}{\partial t} + \nabla \cdot \mathsf{T} = 0, \quad \mathsf{T} = (\beta p + \mathbf{B} \cdot \mathbf{b}) \mathsf{I} - \mathbf{B}\mathbf{b} - \mathbf{b}\mathbf{B}$$

Approximate Schur Complement

$$\mathsf{S}\mathbf{v} = \mathbf{v} + \nabla \cdot \mathsf{T},$$

$$\mathsf{T} \equiv h^2 \theta^2 \left\{ [\mathbf{B} \cdot \nabla \times (\mathbf{v} \times \mathbf{B}) - \gamma \beta \nabla \cdot \mathbf{v}] \mathsf{I} - \mathbf{B} \nabla \times (\mathbf{v} \times \mathbf{B}) - \nabla \times (\mathbf{v} \times \mathbf{B}) \mathbf{B} \right\}$$

Hall MHD Waves

Linearized, Normalized Equations

$$\frac{\partial p}{\partial t} + \gamma \nabla \cdot \mathbf{v} = 0$$

$$\frac{\partial \mathbf{b}}{\partial t} - \nabla \times \left(\mathbf{v} \times \mathbf{B} - d_i \frac{\partial \mathbf{v}}{\partial t} \right) = 0$$

$$\frac{\partial \mathbf{v}}{\partial t} + \nabla \cdot \mathsf{T} = 0, \quad \mathsf{T} = (\beta p + \mathbf{B} \cdot \mathbf{b})\, \mathsf{I} - \mathbf{B}\mathbf{b} - \mathbf{b}\mathbf{B}$$

Approximate Schur Complement

$$\mathsf{S}\mathbf{v} = \mathbf{v} + \nabla \cdot \mathsf{T}$$

$$\begin{aligned} \mathsf{T} \equiv \quad & h^2\theta^2 \left\{ [\mathbf{B} \cdot \nabla \times (\mathbf{v} \times \mathbf{B}) - \gamma\beta \nabla \cdot \mathbf{v}]\, \mathsf{I} - \mathbf{B} \nabla \times (\mathbf{v} \times \mathbf{B}) - \nabla \times (\mathbf{v} \times \mathbf{B})\, \mathbf{B} \right\} \\ & - h\theta d_i \left[(\mathbf{B} \cdot \nabla \times \mathbf{v})\, \mathsf{I} - \mathbf{B}\, (\nabla \times \mathbf{v}) - (\nabla \times \mathbf{v})\, \mathbf{B} \right]. \end{aligned}$$

Static Condensation

- Implicit time step requires linear system solution: **L u = r**.
- Direct solution time grows as n^3.
- Break up large matrix into smaller pieces: Interiors + Interface.
- Small direct solves for interior.
- Interface solve by CG or GMRES, precoditioned with LU or ILU(k) on each processor, with Schwarz overlap between processors.
- Substantially reduces solution time, condition number.

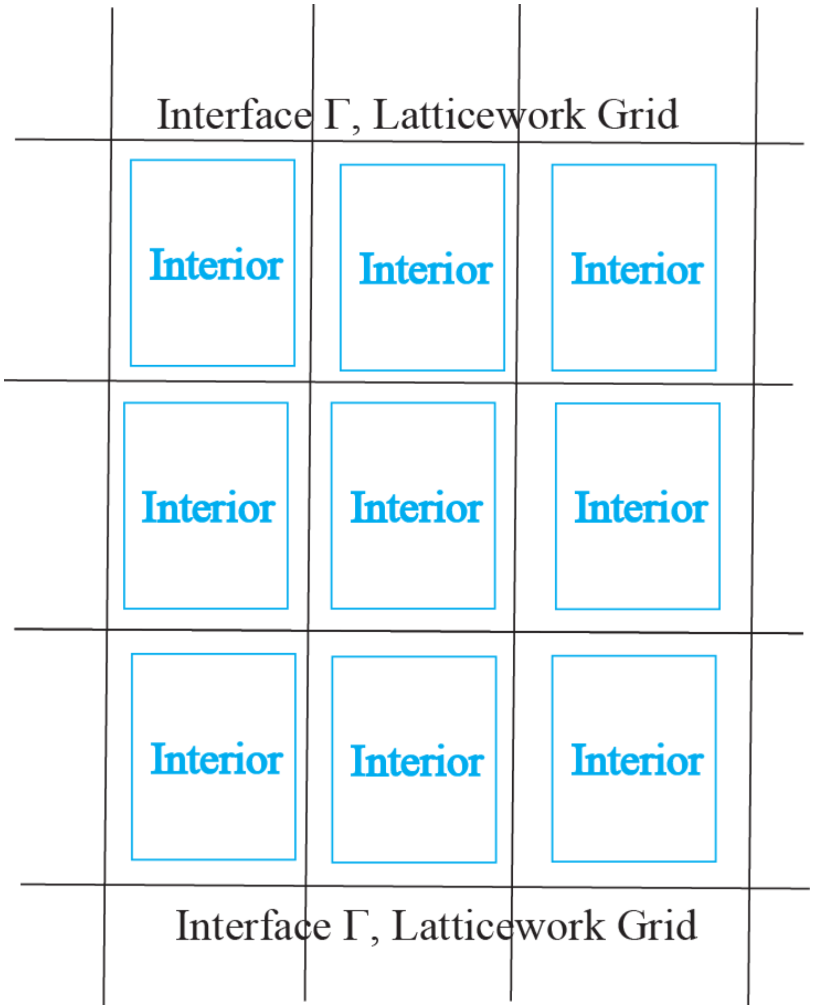

FETI-DP

Finite Element Tearing and Interconnecting, Dual-Primal

- Break up large matrix into three pieces: interior + dual + primal.
- Small direct solves for interior.
- Parallel direct solve for primal points.
- Matrix-free preconditioned GMRES for dual points.
- Primal solve provides information to dual problem about coarse global conditions, providing scalability.
- Interior preconditioner accelerates convergence of dual solve.

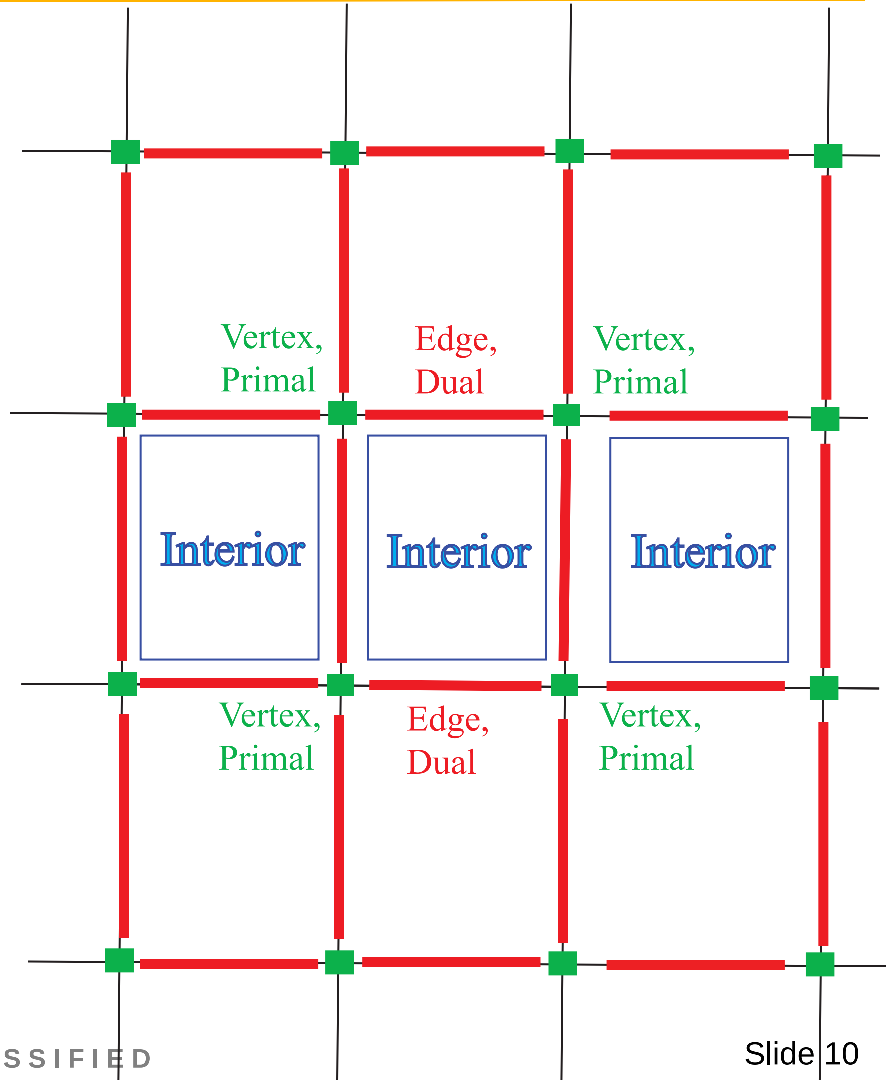

FETI-DP

Finite Element Tearing and Interconnecting, Dual-Primal
Domain decomposition, non-overlapping, Schur complement

Axel Klawonn and Olof B. Widlund,
"Dual-Primal FETI Methods for Linear Elasticity,"
Comm. Pure Appl. Math. **59**, 1523-1572 (2006).

Partition

- I: Interior points, inside each subdomain (grid cell) Ω_i.
- Δ: Dual interface points, continuity imposed by Lagrange multipliers.
- Π: Primal interface points, continuity imposed directly.

Initial Block Matrix Form

$$\mathsf{L}\mathbf{u} = \mathbf{r}, \quad \mathsf{L} = \begin{pmatrix} \mathsf{L}_{II} & \mathsf{L}_{I\Delta} & \mathsf{L}_{I\Pi} \\ \mathsf{L}_{\Delta I} & \mathsf{L}_{\Delta\Delta} & \mathsf{L}_{\Delta\Pi} \\ \mathsf{L}_{\Pi I} & \mathsf{L}_{\Pi\Delta} & \mathsf{L}_{\Pi\Pi} \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} \mathbf{u}_I \\ \mathbf{u}_\Delta \\ \mathbf{u}_\Pi \end{pmatrix}, \quad \mathbf{r} = \begin{pmatrix} \mathbf{r}_I \\ \mathbf{r}_\Delta \\ \mathbf{r}_\Pi \end{pmatrix}$$

Algebraic Reorganization

Local Block Matrices: $\mathbf{I} + \Delta$

$$\mathbf{L}_{BB} = \begin{pmatrix} \mathbf{L}_{II} & \mathbf{L}_{I\Delta} \\ \mathbf{L}_{\Delta I} & \mathbf{L}_{\Delta\Delta} \end{pmatrix}, \quad \mathbf{u}_B = \begin{pmatrix} \mathbf{u}_I \\ \mathbf{u}_\Delta \end{pmatrix}, \quad \mathbf{r}_B = \begin{pmatrix} \mathbf{r}_I \\ \mathbf{r}_\Delta \end{pmatrix}$$

Dual Continuity: Lagrange Multipliers

λ is a vector of Lagrange multipliers used to impose continuity on the dual dependent variables $\mathbf{u}_\Delta$.

$$\mathbf{B} = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_\Delta \end{pmatrix}, \quad \mathbf{B}_\Delta \mathbf{u}_\Delta = 0, \quad \mathbf{L}_{BB}\mathbf{u}_B + \mathbf{L}_{B\Pi}\mathbf{u}_\Pi + \mathbf{B}^T\lambda = \mathbf{r}_B$$

Final Block Matrix Form

$$\mathbf{L} = \begin{pmatrix} \mathbf{L}_{BB} & \mathbf{L}_{B\Pi} & \mathbf{B}^T \\ \mathbf{L}_{\Pi B} & \mathbf{L}_{\Pi\Pi} & \mathbf{0} \\ \mathbf{B} & \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} \mathbf{u}_B \\ \mathbf{u}_\Pi \\ \lambda \end{pmatrix}, \quad \mathbf{r} = \begin{pmatrix} \mathbf{r}_B \\ \mathbf{r}_\Pi \\ \mathbf{0} \end{pmatrix}$$

Solution and Reduction

Solutions for $\mathbf{u}_B$ and $\mathbf{u}_\Pi$

$$\mathbf{u}_B = \mathbf{L}_{BB}^{-1}\left(\mathbf{r}_B - \mathbf{L}_{B\Pi}\mathbf{u}_\Pi - \mathbf{B}^T\lambda\right)$$

$$\mathbf{S}_{\Pi\Pi} \equiv \mathbf{L}_{\Pi\Pi} - \mathbf{L}_{\Pi B}\mathbf{L}_{BB}^{-1}\mathbf{L}_{B\Pi}$$

$$\mathbf{u}_\Pi = \mathbf{S}_{\Pi\Pi}^{-1}\left[\mathbf{r}_\Pi - \mathbf{L}_{\Pi B}\mathbf{L}_{BB}^{-1}\left(\mathbf{r}_B - \mathbf{B}^T\lambda\right)\right]$$

Global Schur Complement Equation for λ

$$\mathbf{F}\lambda = \mathbf{d}$$

$$\mathbf{F} = \mathbf{B}\left(\mathbf{L}_{BB}^{-1} + \mathbf{L}_{BB}^{-1}\mathbf{L}_{B\Pi}\mathbf{S}_{\Pi\Pi}^{-1}\mathbf{L}_{\Pi B}\mathbf{L}_{BB}^{-1}\right)\mathbf{B}^T$$

$$\mathbf{d} = \mathbf{B}\mathbf{L}_{BB}^{-1}\left[\mathbf{r}_B - \mathbf{L}_{B\Pi}\mathbf{S}_{\Pi\Pi}^{-1}\left(\mathbf{r}_\Pi - \mathbf{L}_{\Pi B}\mathbf{L}_{BB}^{-1}\mathbf{r}_B\right)\right]$$

FETI-DP, Variational Formulation
Symmetric Matrix

$$\mathsf{L}\mathbf{u} = \mathbf{r}, \quad \mathsf{B}\mathbf{u} = 0, \quad \mathsf{L}^T = \mathsf{L}$$

$$\mathcal{L} \equiv \frac{1}{2}(\mathbf{u}, \mathsf{L}\mathbf{u}) - (\mathbf{u}, \mathbf{r}) + (\lambda, \mathsf{B}\mathbf{u})$$

$$\frac{\delta \mathcal{L}}{\delta \mathbf{u}} = \mathsf{L}\mathbf{u} - \mathbf{r} + \mathsf{B}^T \lambda = 0$$

$$\frac{\delta \mathcal{L}}{\delta \lambda} = \mathsf{B}\mathbf{u} = 0$$

Solution Strategy

- Small dense block matrices of $\mathbf{L}_{BB}$ solved locally by LAPACK.
- Sparse global, primal matrix $\mathbf{S}_{\Pi\Pi}$ solved by SuperLU.
  - Short-term: redundant on all processors.
  - Medium-term: distributed SuperLU.
  - Long-term: ILU(k)-preconditioned GMRES.
- Global Schur complement matrix $\mathbf{F}$ solved by matrix-free parallel preconditioned GMRES.
- Choose primal interface constraints to provide coarse global problem, ensure scalability. 2D: vertices. 3D: more complicated.
- The scalability of $\mathbf{F}$ is accomplished by the coarse, primal solver. The quality of the preconditioner determines the rate of convergence but not the scalability.

Condition Number and Scalability

Condition Number

$$\mathbf{L}\mathbf{u}_i = \lambda_i \mathbf{u}_i, \quad \kappa(\mathbf{L}) \equiv \frac{\lambda_{\max}}{\lambda_{\min}} \tag{19}$$

Scalability Theorem

A matrix $\mathbf{L}$ is scalable if its condition number, and hence the number of Krylov iterations to convergence, is independent of the number of subdomains. Jan Mandel and Radek Tezaur, "On the convergence of a dual-primal substructuring method," *Numer. Math* **88**, 543-558 (2001). For a symmetric-positive-definite (SPD) matrix, the condition number of the FETI-DP dual matrix is bounded by

$$\kappa(\mathbf{F}) \le C\left[1 + \log^2(H/h)\right], \tag{20}$$

with C constant and H and h characteristic coarse and fine grid spacings.

Weak Scaling Test Problem

- Ideal or Hall MHD waves in a doubly periodic uniform plane.
- 2D **k** vector in computational plane, 3D **B** vector specified by spherical angles about normal to plane. Continuous control of angle θ between **k** and **B.**
- Initialize to pure eigenvector: fast (whistler), shear (kinetic Alfven), or slow wave.
- Unit cell: (knx, kny) full wavelengths.
- Two test cases:
  1. Each processor has one unit cell. Scale up unit cells with nproc. Hold (nx,ny,np) fixed in each unit cell.
  2. One unit cell held fixed, scale up (nx,ny) with nproc. Splits wave length among multiple processors.
- 1 – 64 processors on bassi debug queue.
- Largest test problem size: 16 x 16 wavelengths, 64 processors, 589,824 spatial locations, 3,538,944 variables, 2 large time steps, CFL number ~100, 1 jacobian evaluation, wallclock time ~30 seconds.

FETI-DP Dual Condition Number

MHD Slow Wave, Various k-B Angle θ

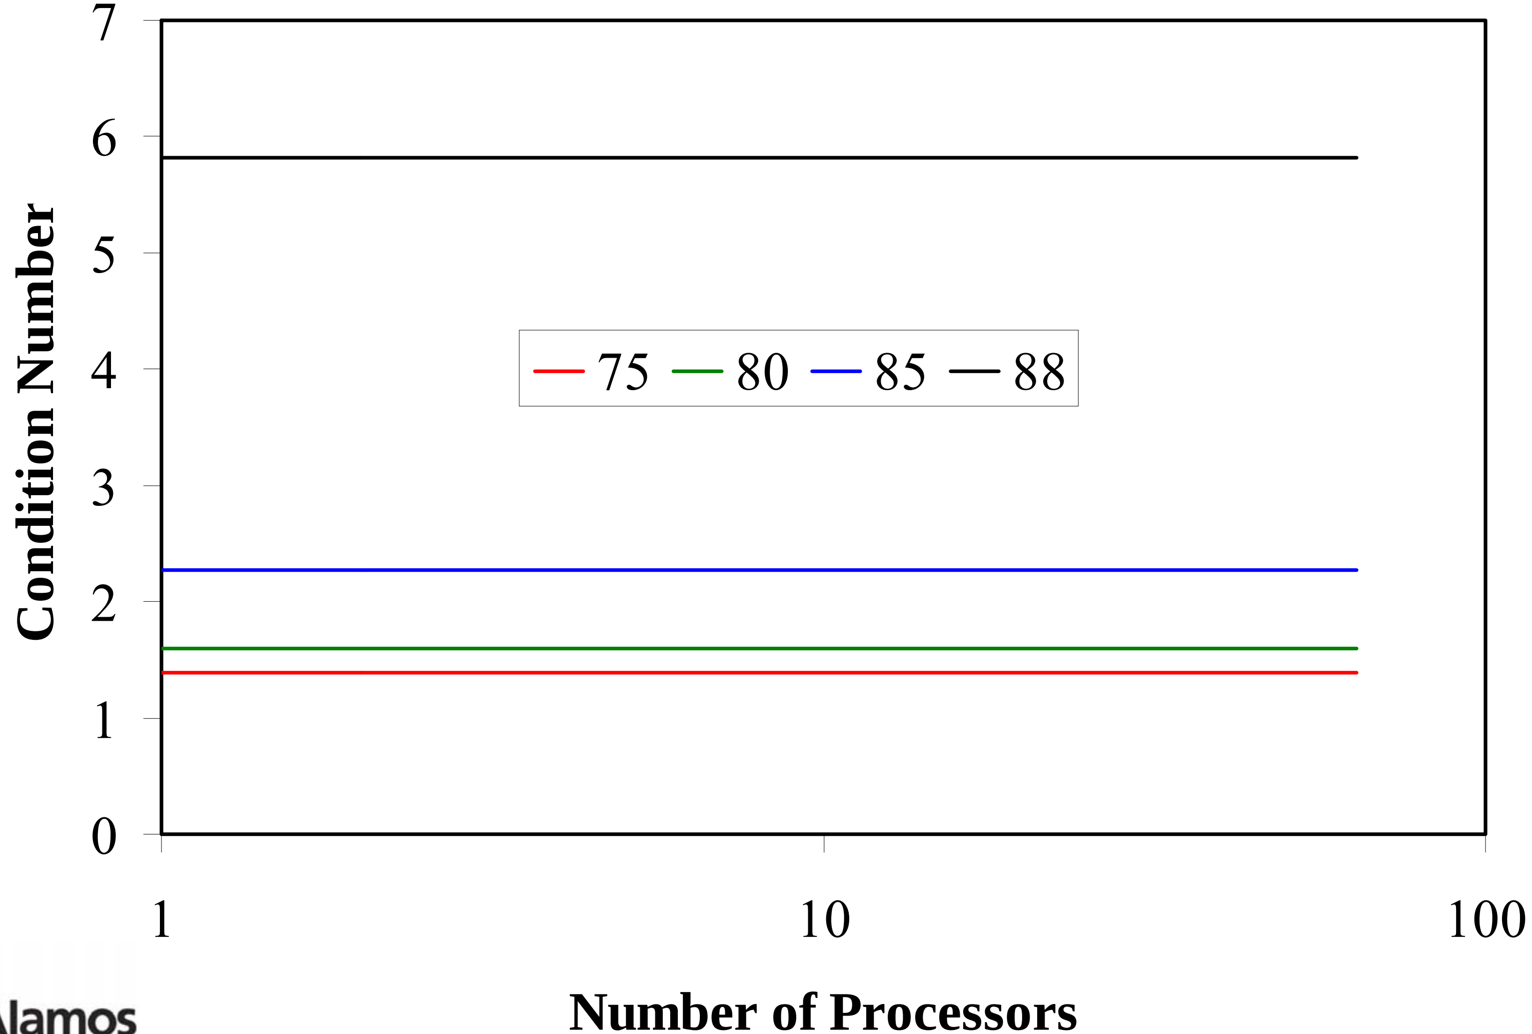

Wallclock Time to Solution

MHD Slow Wave, $\theta = 75^o$, FETI-DP vs. Static Condensation

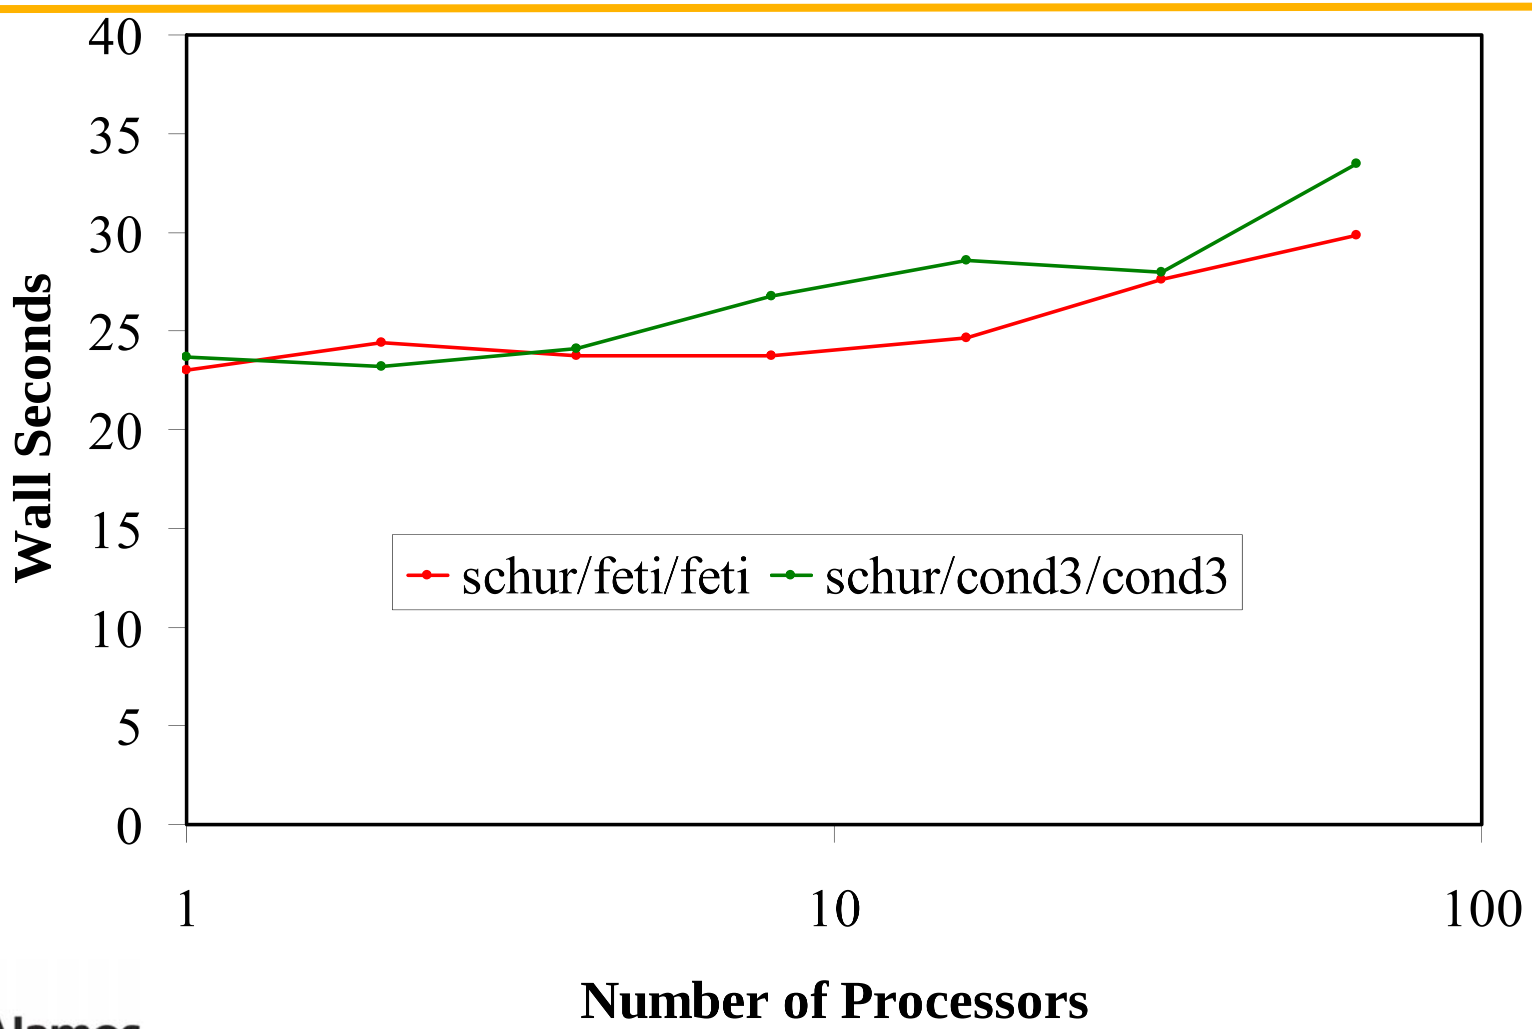

Conclusions

- **Physics-Based Preconditioning**
  - Reduces matrix order requiring solution
  - Improves condition number and diagonal dominance.
  - Similar to time step split, but maintains full nonlinear accuracy.
- **FETI-DP**
  - Provides scalable solver for SPD preconditioning equations, i.e. ideal MHD.
  - Computational results verify analytical scalability theorem.
  - Requires extension to non-SPD problems, such as Hall MHD.
  - Primal solve requires minor modifications to achieve true scalability.
  - 3D primal constraints require research.
- **Static Condensation**
  - Appears to be as scalable as FETI-DP on 1-64 processors.
  - No increase in condition number and time as theta approache 90 degrees.
  - Requires no extension for non-SPD problems.
  - Already implemented for the 3D HiFi spectral element code (Sato).

Future Plans

- Increase number of processors; bassi → franklin.
- Improve parallel implementation, scaling; profiling.
- Test static condensation on Hall MHD Schur complement.
- Investigate extension of FETI-DP to non-SPD problems.
- Continue Schur complement development for dissipative terms, nonuniformity, and nonlinearity.
- Port methods to 3D codes:
  - HiFi
  - M3D
  - NIMROD

Poster Sessions, 9:45 – 12:45 AM Monday

- BP6.00050
  "Development and verification of HiFi -- an adaptive implicit 3D high order finite element code for general multi-fluid applications"
  V. S. Lukin, A. H. Glasser, W. Lowrie, E. Meier, U. Shumlak, M. Sato

- BP6.00051
  "Scalable Parallel Computation for Extended MHD Modeling of Fusion Plasmas"
  Alan H. Glasser