Domain Decomposition: Scalability and Preconditioning A Tutorial

Alan H. Glasser

Los Alamos National Laboratory

Presented at
CEMM Meeting, APS/DPP
Philadelphia, PA, October 29, 2006

Scalability By Domain Decomposition

- 3D extended MHD modeling of magnetically confined fusion plasmas requires petascale computing: 1 petaflop = 10^{15} flops, $\sim 10^4$ procs.
- Efficient petascale computing requires scalable linear systems: condition number independent of grid size, number of processors.
- Domain decomposition is a promising approach to scalability.
  - Schwarz overlapping methods.
  - Non-overlapping methods, domain substructuring, *e.g.* FETI-DP.
- Analytical proofs of scalability for simple systems: Poisson, linear elasticity, Navier-Stokes.
- Empirical studies proposed using existing 2D SEL code for extended MHD.

SEL Code Features

- Flux-source form: simple, general problem setup.
- Spatial discretization:
  - High-order spectral elements, modal basis.
  - Harmonic grid generation, adaptation.
- Time step: fully implicit, 2^{nd}-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

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

Alternative Polynomial Bases

Uniform Nodal Basis

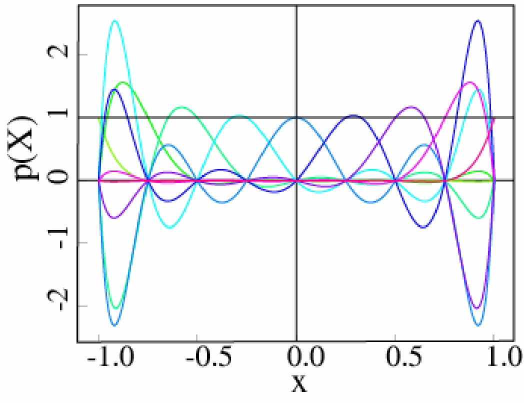

- Lagrange interpolatory polynomials
- Uniformly-spaced nodes
- Diagonally subdominant

Jacobi Nodal Basis

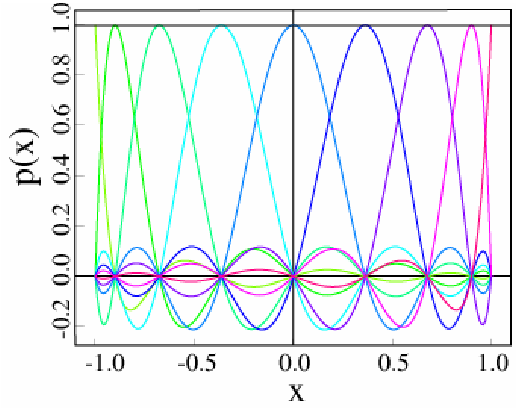

- Lagrange interpolatory polynomials
- Nodes at roots of $(1-x^2)\ P_n^{(0,0)}(x)$
- Diagonally dominant

Spectral (Modal) Basis

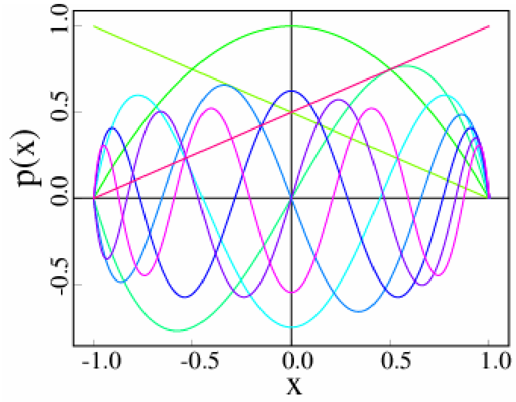

- Jacobi polynomials $(1+x)/2$, $(1-x)/2$, $(1-x^2)\ P_n^{(1,1)}(x)$
- Nearly orthogonal
- Manifest exponential convergence

Implicit Time Discretization: θ-Scheme

$$\mathbf{M}\dot{\mathbf{u}} = \mathbf{r}$$

$$\mathbf{M}\left(\frac{\mathbf{u}^+ - \mathbf{u}^-}{h}\right) = \theta\mathbf{r}^+ + (1-\theta)\mathbf{r}^-$$

$$\mathbf{R}\left(\mathbf{u}^+\right) \equiv \mathbf{M}\left(\mathbf{u}^+ - \mathbf{u}^-\right) - h\left[\theta\mathbf{r}^+ + (1-\theta)\mathbf{r}^-\right] \to 0$$

$$\mathbf{J} \equiv \mathbf{M} - h\theta\left\{\frac{\partial r_i^+}{\partial u_j^+}\right\}$$

$$\mathbf{R}\left(\mathbf{u}^+\right) + \mathbf{J}\delta\mathbf{u}^+ = \mathbf{0}, \quad \delta\mathbf{u}^+ = -\mathbf{J}^{-1}\mathbf{R}\left(\mathbf{u}^+\right), \quad \mathbf{u}^+ \to \mathbf{u}^+ + \delta\mathbf{u}^+$$

- Nonlinear Newton-Krylov iteration.
- Elliptic equations: $\mathbf{M} = 0$.
- Static condensation
- PETSc: GMRES with Schwarz ILU, overlap of 3, fill-in of 5.

Static Condensation

Partition into Subdomains (Grid Cells) Ω_i

I: Interiors
Γ: Interface: (faces) + edges + vertices.

Block Matrix Form

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

Solution for $\mathbf{u}_I$

$$\mathbf{u}_I = \mathbf{L}_{II}^{-1}\left(\mathbf{r}_I - \mathbf{L}_{I\Gamma}\mathbf{u}_\Gamma\right)$$

Schur Complement

$$\mathbf{S} \equiv \mathbf{L}_{\Gamma\Gamma} - \mathbf{L}_{\Gamma I}\mathbf{L}_{II}^{-1}\mathbf{L}_{I\Gamma}, \quad \mathbf{S}\mathbf{u}_\Gamma = \mathbf{r}_\Gamma - \mathbf{L}_{\Gamma I}\mathbf{L}_{II}^{-1}\mathbf{r}_I$$

- $\mathbf{L}_{II}^{-1}$: small local direct solves, LU factorization and back substitution.
- $\mathbf{S}^{-1}$: global solve, preconditioned GMRES.

The Benefits of Static Condensation

nx = number of grid cells in *x* direction
ny = number of grid cells in *y* direction
np = degree of polynomials in *x* and *y*
nqty = number of physical quantities

N = order of global matrix to be solved

Without static condensation: $N = nx\ ny\ nqty\ np^2$
With static condensation: $N = nx\ ny\ nqty\ (2\ np - 1)$

Surface to volume ratio.
Substantial reduction of condition number.

So What's Not To Like? Scalability!

The global Schur complement matrix $\mathbf{S}$ is not scalable. Its condition number, and hence the number of Krylov iterations to convergence, increases with *nx* and *ny*.

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

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

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

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

Dual Continuity: Lagrange Multipliers

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

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

Final Block Matrix Form

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

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

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

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

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

Global Schur Complement Equation for λ

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

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

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

Solution Strategy

- Relatively small dense block matrices of $\mathbf{L}_{BB}$ and sparse matrix $\mathbf{S}_{\Pi\Pi}$ solved by direct LU factorization and back substitution.
- Global Schur complement matrix $\mathbf{F}$ solved by parallel preconditioned Krylov method, *e.g.* GMRES. Requires preconditioner for adequate rate of convergence.
- 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.
- Scalability has been proven analytically for a limited range of simple problems: Poisson, linear elasticity, Navier-Stokes. More general: empirical.

Definitions For Each Subdomain Ω_i

$$\mathbf{B}_{D,\Delta}^{(i)} \equiv \text{ scaled jump matrix}$$

$$\mathbf{R}_{\Gamma\Delta}^{(i)} \equiv \text{ restriction matrix from full interface to dual variables}$$

$$\mathbf{S}_{\varepsilon}^{(i)} \equiv \text{ Schur complement obtained by eliminating interior variables}$$

Preconditioner

$$\mathbf{M}^{-1} = \sum_{i=1}^{n} \mathbf{B}_{D,\Delta}^{(i)} \mathbf{R}_{\Gamma\Delta}^{(i)} \mathbf{S}_{\varepsilon}^{(i)} \mathbf{R}_{\Gamma\Delta}^{(i)T} \mathbf{B}_{D,\Delta}^{(i)T}, \quad \mathbf{M}^{-1}\mathbf{F}\lambda = \mathbf{M}^{-1}\mathbf{d}$$

Condition Number

$$\mathbf{A}\mathbf{u}_i = \lambda_i \mathbf{u}_i, \quad \kappa(\mathbf{A}) \equiv \left| \frac{\lambda_{\max}}{\lambda_{\min}} \right|$$

Scalability

A method is scalable if the condition number of the matrix, and hence the number of Krylov iterations to convergence, is independent of the number of subdomains. $\mathbf{M}^{-1}\mathbf{F}$ has been proven to be scalable for a limited range of physical problems.

Proposed Research Program

- Use existing 2D SEL spectral element code as test bed.
- Implement FETI-DP as a modification of existing static condensation routines.
- Study a progression of extended MHD systems as *nx* and *ny* are increased to determine:
  - Constancy of condition number.
  - Constancy of Krylov iterations required for convergence.
  - Scaling of condition number with parameters.
- Extend spectral element code to 3D.
- Investigate optimal choice of primal constraints for scalability.