Spectral Element Multigrid

A. H. Glasser, PSI Center, University of Washington
V. S. Lukin, Naval Research Laboratory

Presented at the 2010 APS/DPP and CEMM Meetings
Chicago, Illinois, November 7 & 8, 2010

Scalable Parallel Solver Strategies

- **Two-Part Overall Strategy**
  - Physics-based preconditioning (PBP) is used to reduce the order of matrices and make them diagonally dominant (Chacón). Already implemented.
  - Multigrid provides a scalable parallel method for solving diagonally-dominant matrices.
- **Coarsening and Refining Strategies**
  - **Geometric (GMG):** coarser and finer grids. Not suitable for spectral elements.
  - **Algebraic (AMG):** choose largest matrix elements for coarsening. PETSc/Hypre/BoomerAMG. Tested for spectral elements; unsuccessful.
  - **Spectral Element (SEMG):** Specifically designed to exploit spectral elements; uses nest spaces of higher and lower polynomial degrees within each grid cell.
- **Smoothers and Solvers**
  - **Jacobi:** Based on diagonal dominance of nodal basis representation. Used on all fine levels.
  - **Parallel Direct Solve:** Currently used on the coarsest level. Limits scalability.
  - **AMG:** Will be tested as a scalable replacement for Parallel Direct Solve.
- **References**

  - E. M. Ronquist and A. T. Patera, "Spectral Element Multigrid I: Formulation and Numerical Results," J. Sci. Comput. **2**, 4, 389-406 (1987).
  - Y. Madera and R. Muñoz, "Spectral Element Multigrid II: Theoretical Justification," J. Sci. Comput. **3**, 4, 323-353 (1988).

Abstract Multigrid Algorithm

Abstract Elliptic Problem

Hilbert space $\mathcal{H}$, bilinear elliptic form a, linear form g.

Find $v \in \mathcal{H}$ such that $\quad \forall u \in \mathcal{H}, \quad a(u,v) = g(u)$

Nested Finite-Dimensional Subspaces

$\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_j \subset \mathcal{H}$ Simplest case $j = 1, 2$.

Find $v_j \in \mathcal{M}_j$ such that $\quad \forall u \in \mathcal{M}_j, \quad a(u,v_j) = g(u)$

Smoother

$b(u,v) \approx a(u,v)$, but easier to solve, *e.g.* Jacobi smoother, $b \sim \mathrm{Diag}(a)$.

Multigrid V-Cycle

1. $m/2$ smoother iterations. Find $\mathcal{S}\varphi \in \mathcal{M}_2$ such that
$$\forall u \in \mathcal{M}_2, \quad a(u, \mathcal{S}\varphi - \varphi) = g(u) - a(u,\varphi)$$
2. Coarse correction. Find $\bar{\varphi} \in \mathcal{M}_1$ such that
$$\forall u \in \mathcal{M}_1, \quad a(u, \mathcal{S}\varphi - \bar{\varphi}) = g(u) - a(u,\varphi), \quad \mathcal{C}\varphi \equiv \varphi + \bar{\varphi}$$
3. $m/2$ smoother iterations.

Operator Expression

$$u^1 = \mathcal{S}^{m/2}\mathcal{C}\mathcal{S}^{m/2}u^0$$

Convergence Theorem

Generalized Eigenvalue Problem

$$\forall u \in \mathcal{M}_1, \quad a(u, \Psi_i) = \lambda_i b(u, \Psi_i)$$

$$\text{Rescale } b \to \frac{b}{\lambda_{\max}}, \quad 0 < \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_P = 1$$

Nesting Hypothesis

The fine space is the span of all P eigenvectors, $\mathcal{M}_2 = \{\Psi_1, \Psi_2, \ldots, \Psi_P\}$.

The coarse space is the span of the lowest p eigenvectors, $\mathcal{M}_1 = \{\Psi_1, \Psi_2, \ldots, \Psi_p\}$.

Convergence Theorem

Let $e^j \equiv u - u^j$, the error after the jth iteration.

$$a(e^1, e^1) \leq (1 - \lambda_{p+1})^{2m} a(e^0, e^0)$$

Interpretation

The largest "rough" eigenvalues, closest to 1, converge rapidly.

The smallest "smooth" eigenvalues, closest to 0, converge slowly.

The coarse correction eliminates the smooth eigenvalues
by transferring them to a coarse grid and solving exactly.

Coarsening and Refining

Lobatto Nodal Basis Functions

$$u(x) = u_i \alpha_i(x), \quad x \in (-1,1), \quad i = 0, \ldots, n$$

$$\alpha_i(x) \equiv \prod_{j \neq i} \left(\frac{x - \xi_j}{\xi_i - \xi_j} \right), \quad (1 - \xi_i^2) P_n^{(0,0)\prime}(\xi_i) = 0$$

Coarsening and Refining Operations

$$u^m(x) = \sum_{i=0}^{m} u_i^m \alpha_i^m(x), \quad u^n(x) = \sum_{i=0}^{n} u_i^n \alpha_i^n(x), \quad m < n$$

$$C_{ij}^{mn} \equiv \alpha_i(\xi_j^n), \quad u_i^m = C_{ij}^{mn} u_j^n, \quad i = 0, \ldots, m, \quad j = 0, \ldots, n$$

$$\mathbf{C}^{mn} = \{C_{ij}^{mn}\}, \quad \mathbf{u}^m = \mathbf{C}^{mn} \mathbf{u}^n$$

$$\mathbf{R}^{nm} = (\mathbf{C}^{mn})^T, \quad \mathbf{u}^n = \mathbf{R}^{nm} \mathbf{u}^n$$

Lobatto Nodal Basis Functions

Coarse: np = 4

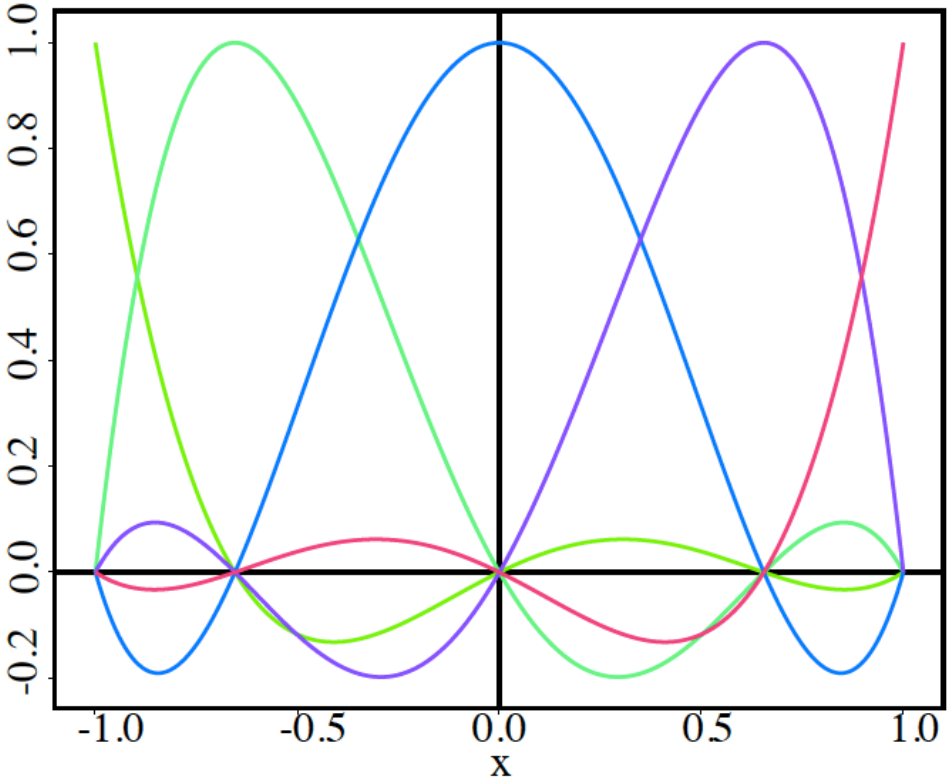

Fine: np = 8

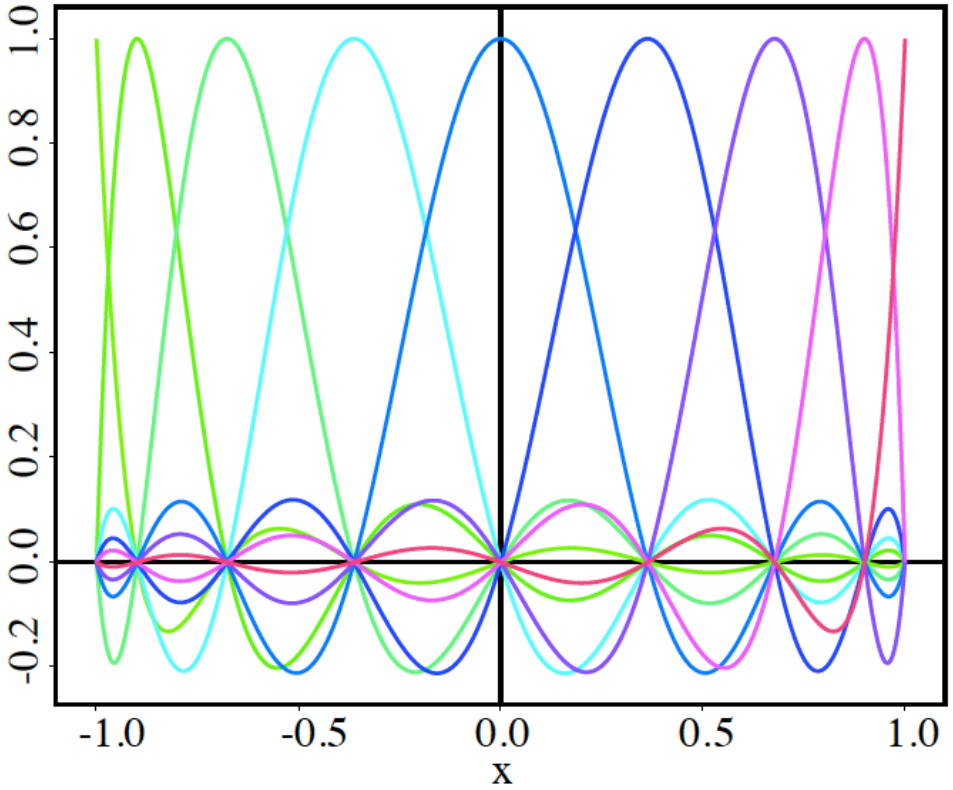

Coarsener: evaluating fine basis functions at coarse nodes.
Refiner is the transpose of the coarsener

1D Stiffness Matrix

Coarse and Fine Stiffness Matrices

$$(u, Lv) \equiv -\int_{-1}^{1} dx u \frac{\partial^2 v}{\partial x^2} = \int_{-1}^{1} dx \frac{\partial u}{\partial x}\frac{\partial v}{\partial x} = u_i A_{ij} v_j$$

$$A^m_{ij} \equiv \int_{-1}^{1} dx \frac{\partial \alpha^m_i}{\partial x}\frac{\partial \alpha^m_j}{\partial x}, \quad i,j = 1,\ldots,m$$

$$A^n_{kl} \equiv \int_{-1}^{1} dx \frac{\partial \alpha^n_k}{\partial x}\frac{\partial \alpha^n_l}{\partial x}, \quad k,l = 1,\ldots,n$$

$m < n$, $m =$ coarse, $n =$ fine.

Coarsening the Stiffness Matrix

Given the fine stiffness matrix,
we can coarsen it with the refining matrix $\mathbf{R}$.

$$u^m_i A^m_{ij} v^m_j = R^{nm}_{ki} u^m_i A^n_{kl} R^{nm}_{lj} v^m_j$$

$$A^m_{ij} = R^{nm}_{ki} A^n_{kl} R^{nm}_{lj}, \quad \mathbf{A}^m = (\mathbf{R}^{nm})^T \mathbf{A}^n \mathbf{R}^{nm}$$

Analytical Test Case

1D Poisson Equation

$$-u''(x) = g(x), \quad u(x) = 0 \text{ at } x = \pm 1$$

Convergence Theorem

Using the Lobatto nodal basis functions of degree N, with the scaled diagonal matrix as the smoother, convergence for one element and 2 levels is given by

$$a(u - u^1, u - u^1) = \mu a(u - u^0, u - u^0)$$

$$\mu = \left[1 - \frac{N+2}{2(N-1)}\right]^{2m}$$

$$\mu < 1 \text{ for } N > 2$$

Implementation of SEMG in the 2D HiFi Code

- SEMG is implemented in a Fortran 95 module, 3300 lines of new code written and tested.
- Uses PETSc library for distributed parallel operation and high-level matrix operations.
- Fortran 95 derived types are defined for multigrid levels, diagonal blocks matrices, scatter contexts between native and PETSc vectors, and a full SEMG problem.
- Coarsening and refining use low-order (np+1) matrices **C** and **R**.
- Scaled diagonal Jacobian smoother uses low-order (nqty) matrix **D**.
- Jacobian scale factor, largest eigenvalue, is computed by power iteration.
- Recursive Fortran 95 subroutine calls itself at successively coarser levels, with a full parallel direct solve on the coarsest level.
- Highly modular code structure, facilitating easy modification, testing, and improvement.

Weak Parallel Scaling Study: Test Case

2D Poisson Equation

$$-\nabla^2 u(x,y) = \rho(x,y)$$

Homogeneous Dirichlet Boundary Conditions

$$u(0,y) = u(1,y) = u(x,0) = u(x,1) = 0$$

Random Source Term and Initial Conditions

$$\rho(x,y) = \sum_{m=1}^{m_{\max}} \sum_{n=1}^{n_{\max}} \rho_{mn} \sin\left(\frac{mx}{\pi}\right) \sin\left(\frac{ny}{\pi}\right)$$

$$u(x,y) = \sum_{m=1}^{m_{\max}} \sum_{n=1}^{n_{\max}} u_{mn} \sin\left(\frac{mx}{\pi}\right) \sin\left(\frac{ny}{\pi}\right)$$

ρ_{mn} and u_{mn} random

Weak Scaling Study

Hold the work per processor fixed.

Successively double n_x, n_y, $m_{\max}$, and $n_{\max}$
and quadruple the number of paralllel cores.

V-cycle: polynomial degree $n_p = 8, 4, 2, 1, 2, 4, 8$

Jacobi iterations per multigrid level $= 3$
Multigrid cycles to 10^{-5} convergence $= 27$

Weak Parallel Scaling Study: Random Solutions

Initial State

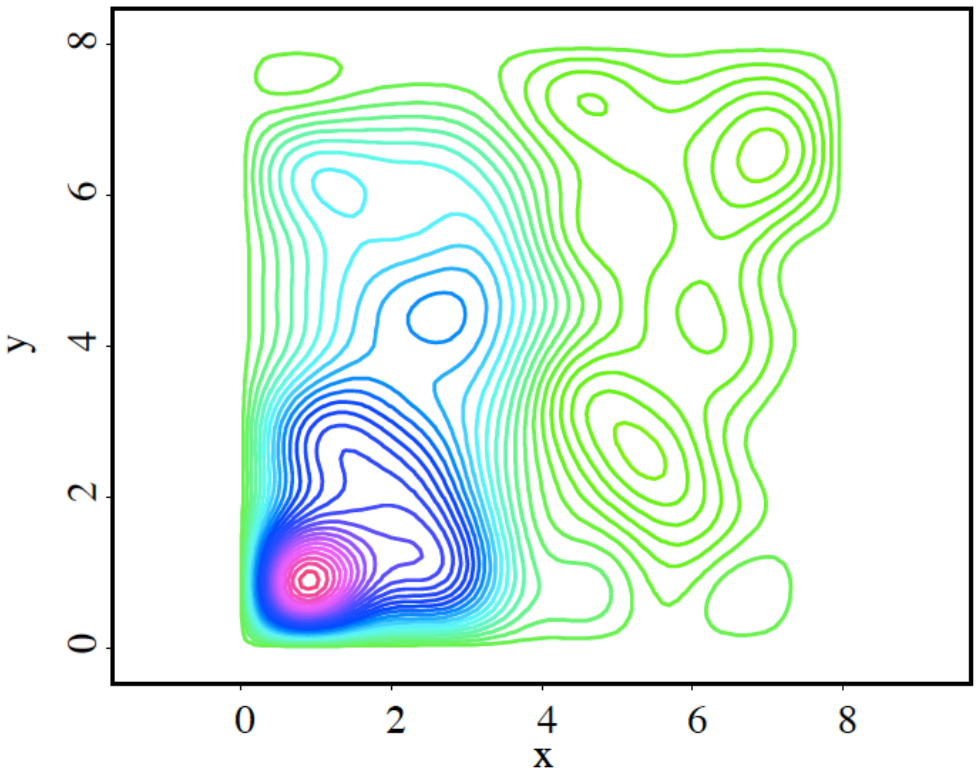

Final State

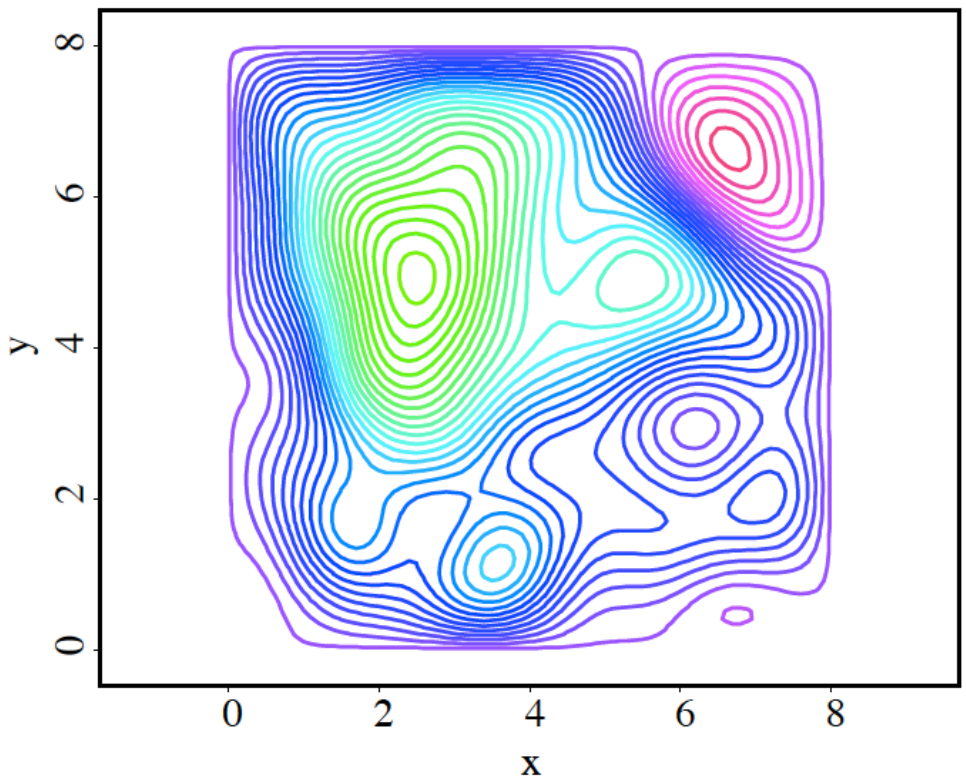

Weak Parallel Scaling Study: Results

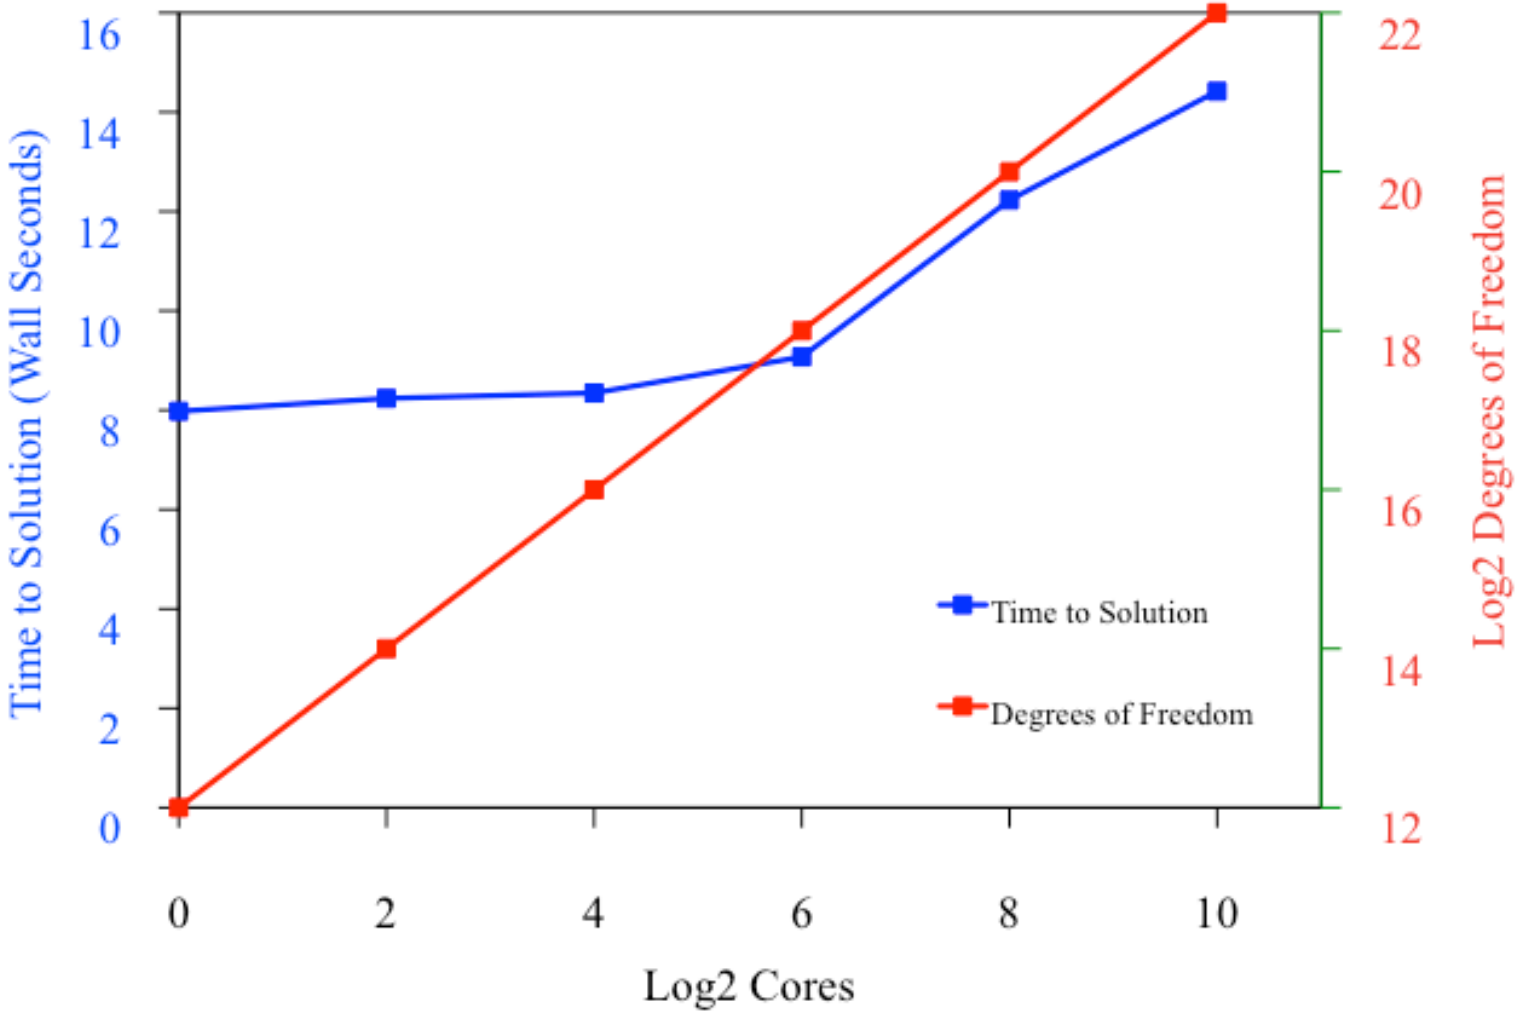

Summary and Future Work

- Formulation and implementation of SEMG are general.
- Tested for scalability on the 2D Poisson equation.
- Improvements to be studied:
  - More scalable alternative to the direct coarse solve, *e.g.* AMG.
  - More general boundary conditions.
- Following this, it will be tested on:
  - Simple linear 2D wave equation.
  - Ideal MHD waves in a periodic plane.
  - GEM challenge problem.
- After completion of 2D development, it will be ported to the 3D HiFi and M3D-C1 codes. The latter will require the development of a different coarsener.