Using Separability to Accelerate PDE Solvers

On the Schur Complement of Nearest Kronecker Product Approximation

Pablo Brubeck

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

APS April Meeting 2018
Columbus, OH April 17, 2018

Outline

Outline

We know how to discretize PDEs, but can we solve them fast?

The goal is to obtain a fast preconditioner for Spectral Galerkin ellipic solvers.

We know how to discretize PDEs, but can we solve them fast?

The goal is to obtain a fast preconditioner for Spectral Galerkin ellipic solvers.

- Applications to steady state problems, eigenmode computation, and implicit time stepping.

We know how to discretize PDEs, but can we solve them fast?

The goal is to obtain a fast preconditioner for Spectral Galerkin ellipic solvers.

- Applications to steady state problems, eigenmode computation, and implicit time stepping.
- Arbitrary geometries, curved quadrilateral elements.

We know how to discretize PDEs, but can we solve them fast?

The goal is to obtain a fast preconditioner for Spectral Galerkin ellipic solvers.

- Applications to steady state problems, eigenmode computation, and implicit time stepping.
- Arbitrary geometries, curved quadrilateral elements.
- High polynomial order (thousands of gridpoints per element, matrix sizes of millions).

Benchmark problem: Newtonian star binary.

- Poisson equation in axisymmetry for the gravitational potential.

Benchmark problem: Newtonian star binary.

- Poisson equation in axisymmetry for the gravitational potential.
- Localized mass density.

Benchmark problem: Newtonian star binary.

- Poisson equation in axisymmetry for the gravitational potential.
- Localized mass density.
- Solution has discontinuous derivatives at the stars' surfaces, single domain methods fail.

Benchmark problem: Newtonian star binary.

- Poisson equation in axisymmetry for the gravitational potential.
- Localized mass density.
- Solution has discontinuous derivatives at the stars' surfaces, single domain methods fail.
- We require a grid that captures the geometry of the problem.

Newtonian star binary, axisymmetric 2D mesh.

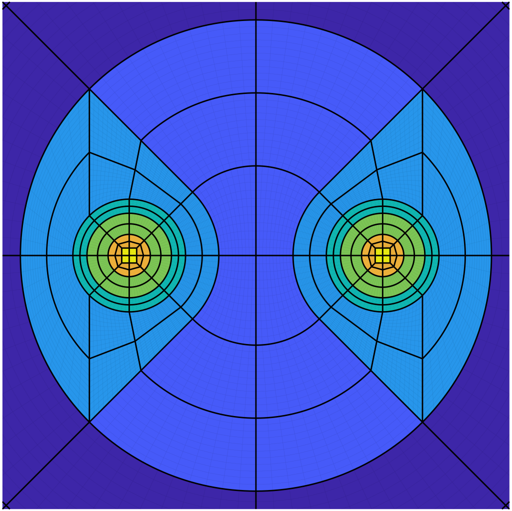

We turn the PDE into an algebraic system of equations.

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

with boundary conditions $\alpha u + \beta \frac{\partial u}{\partial n} = g$ on $\partial\Omega$.

We turn the PDE into an algebraic system of equations.

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

with boundary conditions $\alpha u + \beta \frac{\partial u}{\partial n} = g$ on $\partial\Omega$.

$$A\underline{u} = \underline{f}.$$

We turn the PDE into an algebraic system of equations.

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

with boundary conditions $\alpha u + \beta \frac{\partial u}{\partial n} = g$ on $\partial\Omega$.

$$A\underline{u} = \underline{f}.$$

We make use of a preconditioned iterative solver: approximate A to something that is easier to invert.

The computational domain has regular structure.

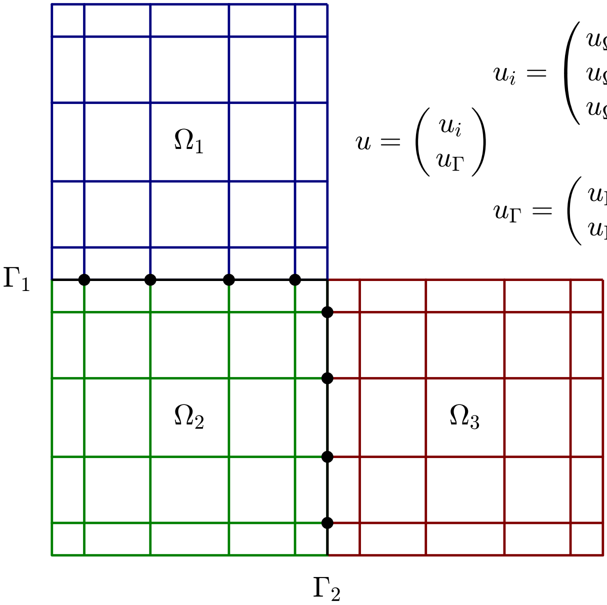

$$u_i = \begin{pmatrix} u_{\Omega,1} \\ u_{\Omega,2} \\ u_{\Omega,3} \end{pmatrix}$$

$$u = \begin{pmatrix} u_i \\ u_\Gamma \end{pmatrix}$$

$$u_\Gamma = \begin{pmatrix} u_{\Gamma,1} \\ u_{\Gamma,2} \end{pmatrix}$$

We can exploit two degrees of separability.

- **Separate dimensions.** Algebraic problems can be solved through separation of variables.

We can exploit two degrees of separability.

- **Separate dimensions.** Algebraic problems can be solved through separation of variables.
- **Separate elements.** Decouple subdomains by first solving for the values at the interface u_Γ.

We get the following recipe:

1. **Separate dimensions.** For each element, approximate the differential operator with one that can be inverted by separation of variables (Nearest Kronecker Product).

We get the following recipe:

1. **Separate dimensions.** For each element, approximate the differential operator with one that can be inverted by separation of variables (Nearest Kronecker Product).
2. **Separate elements.** Solve first for the interface DoFs (Schur Complement), and then solve for the interior DoFs in parallel.

We get the following recipe:

1. **Separate dimensions.** For each element, approximate the differential operator with one that can be inverted by separation of variables (Nearest Kronecker Product).
2. **Separate elements.** Solve first for the interface DoFs (Schur Complement), and then solve for the interior DoFs in parallel.
3. **Rinse and repeat.** Feed this approximation to your favorite iterative solver (Conjugate Gradients, GMRES, etc.).

Outline

Iteration count is independent of problem size.

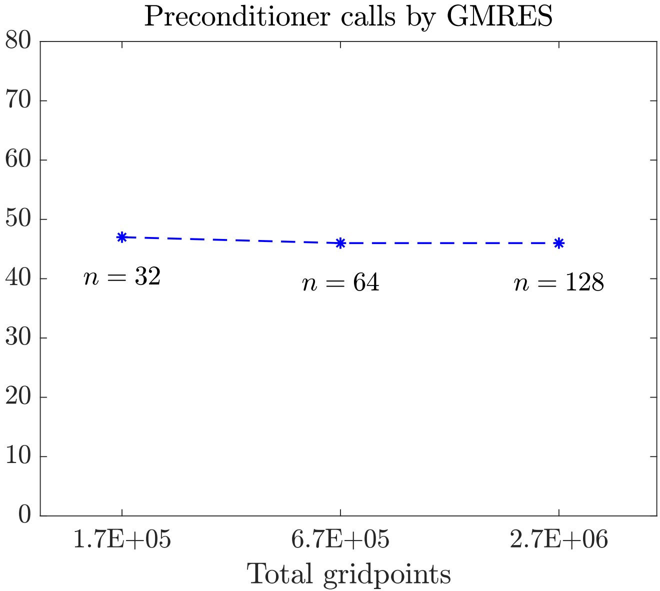

Results

- Explicit computation of the Schur complement of the nearest Kronecker product approximation in $\mathcal{O}(En^3)$ (En^2 grid points).

Results

- Explicit computation of the Schur complement of the nearest Kronecker product approximation in $\mathcal{O}(En^3)$ (En^2 grid points).
- Trivial E-fold parallelization.

Results

- Explicit computation of the Schur complement of the nearest Kronecker product approximation in $\mathcal{O}(En^3)$ (En^2 grid points).
- Trivial E-fold parallelization.
- The preconditioned iterative solver achieves machine precision relative residual.

Results

- Explicit computation of the Schur complement of the nearest Kronecker product approximation in $\mathcal{O}(En^3)$ (En^2 grid points).
- Trivial E-fold parallelization.
- The preconditioned iterative solver achieves machine precision relative residual.
- The iteration count does not scale with the number of elements or polynomial order, it only depends on the equation and the geometry of the domain.

Results

- Explicit computation of the Schur complement of the nearest Kronecker product approximation in $\mathcal{O}(En^3)$ (En^2 grid points).
- Trivial E-fold parallelization.
- The preconditioned iterative solver achieves machine precision relative residual.
- The iteration count does not scale with the number of elements or polynomial order, it only depends on the equation and the geometry of the domain.
- Without preconditioning scaling would have been superlinear.

Outline

Conclusions

- The method presented exploits separability in dimensions and elements.

Conclusions

- The method presented exploits separability in dimensions and elements.
- Explicit computation of the Schur complement of the nearest Kronecker product approximation in $\mathcal{O}(En^3)$ (En^2 gridpoints).

Conclusions

- The method presented exploits separability in dimensions and elements.
- Explicit computation of the Schur complement of the nearest Kronecker product approximation in $\mathcal{O}(En^3)$ (En^2 gridpoints).
- The preconditioned iterative solver achieves machine precision relative residual.

Conclusions

- The method presented exploits separability in dimensions and elements.
- Explicit computation of the Schur complement of the nearest Kronecker product approximation in $\mathcal{O}(En^3)$ (En^2 gridpoints).
- The preconditioned iterative solver achieves machine precision relative residual.
- Iteration count does not scale with problem size.

Future work

- Generalization to 3D.

Future work

- Generalization to 3D.
- Large-scale (HPC) parallel implementation.

Future work

- Generalization to 3D.
- Large-scale (HPC) parallel implementation.
- Applications to implicit time-stepping.

Future work

- Generalization to 3D.
- Large-scale (HPC) parallel implementation.
- Applications to implicit time-stepping.
- Extension to nonlinear problems (Homotopy Analysis Method).

NCSA Gravity Group

1. Daniel George: Deep Learning for Real-time Gravitational Wave Detection and Parameter Estimation: Results with Advanced LIGO Data
2. Adam Rebei: Influence of higher-order waveform multipoles for the detection of eccentric binary black hole mergers
3. Roland Haas: Assessing confidence in numerical relativity waveforms of binary neutron star mergers
4. Hongyu Shen: Glitch Classification and Clustering for LIGO with Deep Transfer Learning
5. Eliu Huerta: Detection and characterization of eccentric compact binary coalescence at the interface of numerical relativity, analytical relativity and machine learning
6. Hongyu Shen: Denoising Gravitational Waves using Deep Learning with Recurrent Denoising Autoencoders (poster)
7. Roland Haas: BOSS-LDG using Blue Waters for LIGO data analysis (poster)
8. Vedant Puri: Scheduled Relaxation Jacobi Method for Initial Data Problems
9. Shawn Rosofsky: Study of f-mode Oscillations in Numerical Relativity Simulations of Perturbed Neutron Stars and Highly Eccentric Binary Neutron Star Mergers
10. Pablo Brubeck: On the Schur complement of the nearest Kronecker product preconditioner for elliptic boundary value problems
11. Haris Markakis: Helmholtz's third theorem in numerical general relativity
12. Miguel Holgado: Pulsar Timing Constraints on the Fermi Massive Black-Hole Binary Blazar Population

References

- W. Pazner and P.-O. Persson, "Approximate tensor-product preconditioners for very high order discontinuous galerkin methods," *Journal of Computational Physics*, vol. 354, pp. 344–369, feb 2018.
- C. F. Van Loan and N. Pitsianis, "Approximation with kronecker products," in *Linear algebra for large scale and real-time applications*, pp. 293–314, Springer, 1993.