Computer science-based approaches for the design of advanced materials

Krylov Space Approximations for Materials Science Problems

Principal Investors: Eric de Sturler[1], Glaucio Paulino[2], David M. Ceperley[3,4], Jeongnim Kim[4,5]
Students: Zhen Cheng[1], Greg Mackey[6], Mike Parks[6], Chris Siefert[1], Shun Wang[1]
Departments of [1]Computer Science, [2]Civil Engineering, and [3]Physics; [4]National Center for Supercomputing Applications; [5]Materials Computation Center; [6]Sandia National Laboratory

Overview:

Approximations from Search Spaces

For many large problems, such as predicting materials behavior, we build a search space, project the problem onto the search space, and solve a small problem. The large problem is then solved as a sequence of small problems.

As examples, we do this for:

- Linear systems and least squares
- Eigenvalue problems
- Systems of ODEs
- Approximations to functions of matrices

The search space we consider here is a Krylov space, a typical element in the search space is a polynomial in a matrix times a vector.

- Cheap in memory
- Cheap in work (if matrix-vector product is cheap, often linear)
- Stop early with adequate accuracy

Many problems lead to long sequences of slowly changing matrices. Based on the underlying mathematics, we select a subspace of previous search spaces and then adapt and reuse ("recycle") this space to improve the convergence or the approximation for subsequent matrices in the sequence.

To make recycling work well, we are concerned with effects of small changes in the matrix: perturbations of solutions, perturbations of invariant subspaces (or eigenvectors), and perturbations of other relevant values, vectors, or spaces.

Our primary objectives are:

- determining for a given problem whether recycling is appropriate and efficient. For example, when simulating crack propagation, the linear solver is faster if we remove low-frequency modes from the equation. The low frequency modes do not change for a tiny propagation of the crack, and are reused for the next matrix. Therefore, the next linear system can be solved faster.
- that the solver learns fast, to cope with relevant system changes. For example, over sufficient propagation of the crack, the low-frequency modes change drastically, but the solver+recycling method is able to cope with the rate of change.

Findings

This idea of "recycling" works well for many applications that deal with a sequence of problems. We have developed new, faster methods for:

- solving linear systems
- solving sequences of linear systems
- estimating determinants of large systems in O(N)

All of this is achieved by recycling over Krylov substeps: carrying out iterative methods that make intelligent use of the results of earlier steps.

Scalability: Current methods for Quantum Monte Carlo have cubic scaling. The methods we are considering appear to have linear or near linear scaling.

Easy to parallelize: Our methods are built on well-recognized, standard, computational components.

Publications arising from this research

- Misha Kilmer and Eric de Sturler, "Recycling Subspace Information for Diffuse Optical Tomography", accepted for publication in *SIAM Journal on Scientific Computing*, July 2005.
- Michael L. Parks, Eric de Sturler, Greg Mackey, Duane D. Johnson, and Spandan Maiti, "Recycling Krylov Subspaces for Sequences of Linear Systems", Tech. Report UIUC DCS-R-2004-2421 (www.cse.uiuc.edu/~sturler/Public/KrylovReuse.ps) and UILU-ENG-2004-1722, March 2004; revision submitted to *SIAM Journal on Scientific Computing*.
- "Fast Iterative Solvers for Topology Optimization", Shun Wang, Eric de Sturler, Glaucio Paulino, to be submitted to *International Journal for Numerical Methods in Engineering* (IJNME).

Collaborations

Applications for this research include crack propagation, fatigue and fracture, diffuse optical tomography, structural topology optimization and design of advanced materials, materials science, Quantum Monte Carlo and variants, electromagnetic wave propagation, and computational fluid dynamics. The research below is funded by the National Science Foundation and the Department of Energy.

Two example applications described on this poster are:

- Structural topology optimization; design of advanced materials (Glaucio Paulino, Civil Eng, UIUC)
- QMC for materials (David Ceperley, Physics UIUC; Jeongnim Kim, MCC/NCSA, UIUC)

Other active research collaborations include:

- Approximating Green's functions in KKR methods (Duane D. Johnson, Matls Sci, UIUC; Andrei Smirnov, ORNL)
- Fracture/crack propagation (Philippe Geubelle, Aero Eng, UIUC)
- Tomography (Misha Kilmer, Tufts)
- Simulation of fracture in disordered materials (Phani Nukala, ORNL)

Application:

Quantum Monte Carlo for Materials

Co-Principal Investigators: David Ceperley (Physics, UIUC) and Eric de Sturler (CS, UIUC)

We are participating in a four-year grant to investigate properties of deep-earth materials, such as ferrite oxides. At depths greater than 100 kilometers, these materials are difficult to evaluate in standard laboratory experiments. One example of where models of the behavior of deep-earth materials could be of use is in understanding the movement of tectonic plates.

In current methods, exactly computing determinants typically has $O(N)^3$ cost. Rather than exactly calculating determinants, our approach is to use subspaces to estimate determinants sufficiently accurately and cheaply, at O(N) cost.

First approach: Using methods of Golub and Bai, based on two steps:

$$u^T f(A)\, v = u^T Q\, f(A)\, Q^T v = \sum_i f(\lambda_i)\, \tilde{u}_i\, \tilde{v}_i = \int_a^b f(\lambda)\, d\mu(\lambda)$$

$$u^T f(A)\, v \approx u^T V_m\, f(T_m)\, V_m^T v$$

Approximating bilinear or quadratic forms over a Krylov space is equivalent to Gaussian quadrature. If we combine the Gaussian quadrature with stochastic techniques, we can estimate determinants, instead of exactly calculating them. This approach scales well but converges slowly, and so is practical only for very large problems.

Current approach: Rather than estimating determinants, we estimate the ratio of determinants of two successive steps, which determines if we accept the next step.

Both our first and current approaches have linear scaling, but the current approach already looks competitive for a small number of (several hundred) electrons.

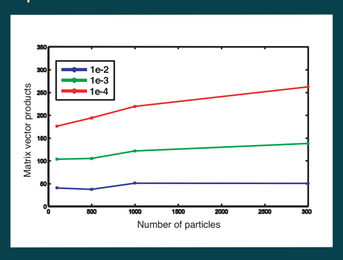

Accuracy in the ratio of the determinants
Image shows number of matrix-vector products required for (approximate) accuracy in MC steps for varying number of particles. The middle line (green) indicates the desired accuracy of 3 digits, and exhibits a very modest increase going from 100 to 3000 particles. Because of theoretical considerations, we believe there is an upper bound for the number of matrix-vector products independent of the number of particles.

Significant results: The stochastic component has been eliminated. The resulting gain in speed may enable computations which are currently performed on supercomputers to be accomplished on a workstation.

This collaboration of mathematical science with geoscience includes Carnegie Institute in Washington, the Ohio State University, Cornell University, and North Carolina State University under NSF EAR Award 05-30643.

Application:

Topology optimization for designing advanced materials

Co-Principal Investigators: Glaucio Paulino (Civil Eng, UIUC) and Eric de Sturler (CS, UIUC)

Computer science plays a key role in modelling and creating materials that have very special properties and behavior, for example, a material with negative thermal expansion coefficient that expands upon cooling.

We have a computed an optimal design for a functional graded material. Due to the functionally graded material properties of each element, pictured at right, we achieve macroscopically a negative Poisson ratio.

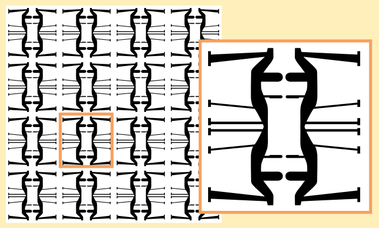

Advanced materials
Left: We have computed an optimal design from functional graded material with a negative Poisson ratio.
Right: a single element of the material.

An element's behavior and density can be represented by equations, and by solving discretized versions of these equations we can calculate total deformation and model and optimize an element's behavior and interaction with neighboring elements, and thus the behavior of the material.

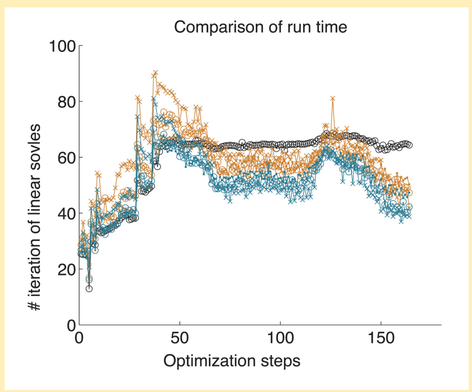

Speed gains
The graph shows the number of linear solver iterations per optimization step. The black line represents a standard method, but with scaling and preconditioning. Our method uses recycling and preconditioning (shown in brown and blue) and is 40-60% faster. The initial optimization steps correspond to large changes in design variables.

Significant results: With one student, we are currently optimizing structures on a PC using an order of magnitude more unknowns than other groups. With another student, we are advancing these computational techniques and materials models to design advanced micro-materials. Other researchers are developing techniques to actually construct such materials.

Future work: The next step is to couple these models and add atomistic/electronic structure calculations to compute appropriate materials models, and perform large multiscale simulations.

Optimization demonstration: Designing a truss-like structure

The figure below illustrates the major steps of the topology optimization algorithm. The change in shape represents changes in design variables.

In this example, the goal is to construct the structure for the strongest bridge made from a given amount of material. The design process is to divide the block into a mesh with many small elements and optimize the local densities or other design variables for the elements. When taking into account the amount of energy taken up by structure itself, the finite element analysis will develop "automatically" into truss-like structures.

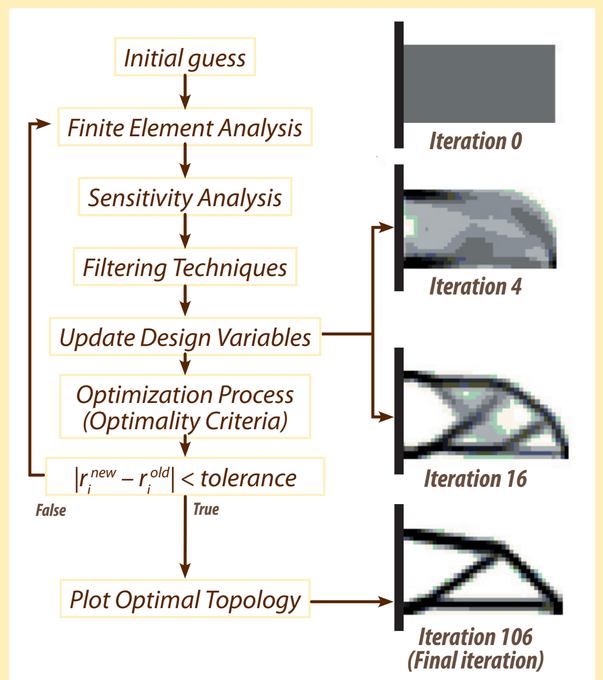

Materials Computation Center
University of Illinois at Urbana-Champaign • Supported by NSF DMR 03-25939

As the Computational Materials Science discipline affects all fields of Science and Engineering, the Materials Computation Center (MCC) is actively developing powerful, leading-edge tools to analyze and predict the properties of materials. MCC provides an intellectual and interactive environment for students, teachers, and researchers focused on world-class, multidisciplinary education and research in Computational Materials Science.

The MCC is headquartered and supported by facilities and services of the Frederick Seitz Materials Research Laboratory. www.mcc.uiuc.edu