

**ABSTRACTS OF THE
FOURTEENTH COPPER MOUNTAIN
CONFERENCE
ON
MULTIGRID METHODS**

**COPPER MOUNTAIN, COLORADO
March 22 - March 27, 2009**

ORGANIZED BY
FRONT RANGE SCIENTIFIC COMPUTATIONS, INC.

CO-ORGANIZED BY
**CENTER FOR APPLIED SCIENTIFIC COMPUTING,
LAWRENCE LIVERMORE NATIONAL LABORATORY
UNIVERSITY OF COLORADO**

SPONSORED BY
**DEPARTMENT OF ENERGY
NATIONAL SCIENCE FOUNDATION
IBM
LOS ALAMOS NATIONAL LABORATORY
LAWRENCE LIVERMORE NATIONAL LABORATORY
SANDIA NATIONAL LABORATORIES**

IN COOPERATION WITH
SIAM

CONFERENCE CHAIR

Van Emden Henson, *Lawrence Livermore National Laboratory*

PROGRAM CO-CHAIRS

Ulrich Rde, *University of Erlangen*

Irad Yavneh, *Technion*

PROGRAM COMMITTEE

Susanne Brenner, *University of South Carolina*

Joel Dendy, *Los Alamos National Laboratory*

Craig Douglas, *University of Wyoming*

Robert D. Falgout, *Lawrence Livermore National Laboratory*

Jim Jones, *Florida Institute of Technology*

Kirk Jordan, *IBM, Watson Research Center*

Tom Manteuffel, *University of Colorado*

Steve McCormick, *University of Colorado*

David Moulton, *Los Alamos National Laboratory*

Kees Oosterlee, *Delft University of Technology*

Joseph Pasciak, *Texas A & M*

John Ruge, *University of Colorado*

Klaus Stben, *Fraunhofer*

Ulrike Yang, *Lawrence Livermore National Laboratory*

Ludmil Zikatanov, *Penn State*

CIRCUS CHAIR

Craig Douglas, *University of Wyoming*

Fourteenth Copper Mountain Conference on Multigrid Methods

March 22 - March 27, 2009

Welcome to the 14th Copper Mountain Conference on Multigrid Methods. This publication is divided into two parts. The first part consists of the Student Paper Competition section. The first three entries of the Student Paper Competition section are the winners of the student competition. The rest of the competition entries are in alphabetical order by speaker's surname. The second part contains regular program abstracts in alphabetical order by speaker's surname.

The entire student papers submitted for the competition can be viewed online at:

<http://grandmaster.colorado.edu/~copper/2009/StudentCompetition>

1. Student Paper Competition

The first three entries are the winners of the student competition; the rest of the competition entries are in alphabetical order by speaker's surname.

Killian Miller
Algebraic Multigrid for Markov Chains

Department of Applied Mathematics
University of Waterloo
Waterloo, Ontario
Canada
k7miller@uwaterloo.ca

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/KMiller_2009.pdf

An algebraic multigrid (AMG) method is presented for the calculation of the stationary probability vector of an irreducible Markov chain. A modified AMG interpolation formula is designed to produce a nonnegative interpolation operator with unit row sums. It is shown how the adoption of a lumping technique maintains the irreducible singular M-matrix character of the coarse-level operators on all levels. Together, these properties are sufficient to guarantee the well-posedness of the algorithm. Numerical results show how it leads to nearly optimal multigrid efficiency for a representative set of test problems.

Eran Treister
Square & Stretch Multigrid for Stochastic Matrix Eigenproblems

Technion – Israel Institute of Technology
Haifa, Israel
eran@tx.technion.ac.il
Irad Yavneh

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/ETreister_2009.pdf

A novel multigrid algorithm for computing the principal eigenvector of column stochastic matrices is developed. The method is based on the *Exact Interpolation Scheme* multigrid approach of Brandt and Ron [1] whereby the prolongation is adapted to yield a better and better coarse representation of the sought eigenvector. The main novelty of the present approach is in the squaring of the stochastic matrix—followed by a stretching of its spectrum—just prior to the coarse-grid correction process. This procedure is shown to yield good convergence properties, even though a cheap and simple aggregation is used for the restriction and prolongation matrices, which is important for maintaining competitive computational costs. A further contribution of this paper is a novel bottom-up procedure for defining the coarse-grid aggregates.

REFERENCES

- [1] A. BRANDT AND D. RON, *Multigrid solvers and multilevel optimization strategies*, In *Multilevel Optimization and VLSICAD*, J. Cong and J. R. Shinnerl, Eds., Kluwer (Boston), (2003), pp. 1–69.

Liuqiang Zhong
**Optimal Multilevel and Adaptive Finite Element Methods for
Time-harmonic Maxwell Equations**

School of Mathematics and Computational Sciences
Xiangtan University
China
zhonglq@xtu.edu.cn

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/LZhong_2009.pdf

Three related results in this paper will be presented for the finite element discretization of time-harmonic Maxwell equations: a two-grid method, L^2 error estimates and convergence and complexity of adaptive finite element methods. Numerical experiments are carried out to justify the optimality of these results.

Ye Chen
Multigrid methods for Monge-Ampere Equations

Department of Mathematics
Clarkson University
Campus Box 5815
Potsdam, NY 13699-5815
chenye@clarkson.edu
Scott R. Fulton

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/YChen_2009.pdf

Best compact discretizations of a Monge-Ampere equation are found. It is verified that there is no fourth-order compact discretization of the Monge-Ampere equation. Multigrid methods combined with τ -extrapolation can solve the Monge-Ampere equation to fourth-order accuracy. The balanced vortex model is solved 4774 times faster by a multigrid algorithm than a single grid method.

Jintao Cui
Multigrid Methods for Symmetric Discontinuous Galerkin Methods on Graded Meshes

Department of Mathematics
Louisiana State University
Baton Rouge, LA 70803
cui@math.lsu.edu

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/JCui_2009.pdf

In this paper we discuss multigrid solvers for systems resulting from the discretization of second order elliptic boundary value problems by a class of symmetric, stable and consistent discontinuous Galerkin (DG) methods on graded meshes. Quasi-optimal error estimates in both the energy norm and the L_2 norm for this class of DG methods are derived and uniform convergence of the W -cycle multigrid algorithm for the resulting discrete problem is proved. We will present both theoretical and numerical results. This is joint work with Susanne C. Brenner, Thirupathi Gudi and Li-yeng Sung.

Steven Hamilton
New Multigrid Smoothers for the Oseen Problem

Emory University
Mathematics & Computer Science
Atlanta, GA 30322
`sphamil@emory.edu`
Michele Benzi
Eldad Haber

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/SHamilton_2009.pdf

We investigate the performance of smoothers based on the Hermitian/skew-Hermitian (HSS) and augmented Lagrangian (AL) splittings applied to MAC discretizations of the Oseen problem. Both steady and unsteady flows are considered. Local Fourier analysis and numerical experiments on a 2-D lid-driven cavity problem indicate that the proposed smoothers result in h -independent convergence and are fairly robust with respect to the Reynolds number.

Karsten Kahl
A least squares based variational AMG solver for HPD systems

Department of Mathematics
University of Wuppertal
Germany
`kkahl@math.uni-wuppertal.de`

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/KKahl_2009.pdf

This work concerns the development of an efficient algebraic multigrid (AMG) solver for the linear systems arising in lattice simulations of quantum chromodynamics (QCD). We present some first ideas and tests for developing an adaptive algebraic multigrid method for these systems based on Brandt's Bootstrap AMG (BAMG) framework. In our proposed solver, we use a variational adaptive setup algorithm to recursively construct the coarse spaces of the AMG hierarchy, with interpolation built using Brandt's notions of Compatible Relaxation, algebraic distances, and Least Squares fitting of relaxed vectors. In addition, in the resulting solver we use an optimal over-weighting in energy of the coarse-grid correction. We demonstrate the efficacy of these approaches for a set of Gauge Laplacian systems, involving hermitian positive definite matrices with highly disordered coefficients, as encountered in Lattice QCD. We conclude with a brief discussion of future work.

Artem Napov
**V-cycle multigrid convergence analysis: comparison of bounds and
relation with the two-grid convergence factor**

Service de Métrologie Nucléaire
Université Libre de Bruxelles (C.P. 165/84)
50 Av. F.D.Roosevelt
B-1050 Brussels
Belgium
anapov@ulb.ac.be
Yvan Notay

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/ANapov_2009.pdf

We consider multigrid methods with V-cycle for symmetric positive definite linear systems. We consider bounds on the convergence factor that are characterized by a constant which is the maximum over all levels of an expression involving only two consecutive levels. More particularly, we consider the classical bound by Hackbusch [1], a bound by McCormick [2], and a bound obtained by applying the successive subspace correction convergence theory with so-called a -orthogonal decomposition [3,4]. We show that the constants in these bounds are closely related, and hence that these analyses are equivalent from the qualitative point of view, whereas McCormick bound is in fact the best one. We also show some relation with the two-grid convergence factor that helps to understand when an optimal two-grid method leads to a multigrid method which is optimal with V-cycle. Moreover, it turns out that when Fourier analysis can be applied to estimate the two-grid convergence factor, one can derive with little additional effort a bound on the convergence factor for the multigrid method with V-cycle. Considering a typical example, we further show that this bound can provide a satisfactorily sharp estimate of the actual multigrid convergence speed.

REFERENCES

- [1] WOLFGAN HACKBUSCH, *Multi-Grid Methods and Applications*, Springer-Verlag (1985) 164-167.
- [2] STEVE F. MCCORMICK, *Multigrid methods for variational problems: general theory for the V-cycle*, SIAM J. Numer. Anal. (1985) 634-643
- [3] JINCHAO XU, *Iterative Methods by space decomposition and subspace correction*, SIAM Review (1992) 581-613.
- [4] HARRY YSERENTANT, *Old and new convergence proofs for multigrid methods*, Acta Numerica (1993) 285-326.

Daniel Ritter
**A Fast Adaptive Composite Grid Algorithm for Solving the Free – Space
Poisson Problem on the Cell Broadband Engine**

Department of Computer Science 10
University of Erlangen–Nuremberg
Germany daniel.ritter@informatik.uni-erlangen.de

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/DRitter_2009.pdf

In the field of molecular dynamics Poisson’s equation with open boundary conditions plays a major role. One issue that arises when this equation is solved numerically is the infinite size of the domain. This prevents a direct solution so that other concepts have to be considered. Within this paper a method is discussed that uses hierarchical coarsened grids to overcome this problem. Special attention has to be paid to the discretization at the grid interfaces. A finite volume approach has been used for that. The resulting set of linear equations is solved using a fast adaptive composite grid algorithm.

Emphasis is put on the implementation of the method on the STI Cell Broadband Engine, a modern multi core processor, that is powerful in floating point operations and memory bandwidth. Code optimization techniques are applied as well as parallelization of the code to get maximum performance on this processor. For validation of the performance test runs are executed and the runtime is analyzed in detail.

Carmen Rodrigo
Vector Fourier Analysis on Triangular Grids for Planar Elasticity

Department of Applied Mathematics
University of Zaragoza
Spain
carmenr@unizar.es

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/CRodrigo_2009.pdf

In this paper Local Fourier Analysis (LFA) for multigrid methods on triangular grids is extended to the case of systems of PDEs. In particular, it is performed for the problem of planar elasticity, although its application to other systems is straightforward. Analogously to the scalar case, this analysis is based on an expression of the Fourier transform in new coordinate systems, both in space and in frequency variables, associated with reciprocal bases. LFA is particularly valuable for systems of PDEs, since it is often much more difficult to identify the correct multigrid components than for a scalar problem. For the discrete elasticity operator obtained with linear finite elements, different collective smoothers like three-color smoother and some zebra-type smoothers are analyzed. LFA results for these smoothers are presented.

Alex Thekale
Optimizing the number of multigrid cycles in the full multigrid algorithm

Applied Mathematics II
Department of Mathematics
University of Erlangen-Nuremberg
Martensstr. 3
91058 Erlangen
Germany
`Alexander.Thekale@am.uni-erlangen.de`

For complete paper, please visit
<http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/AThekale.2009.pdf>

Multigrid methods are among the most efficient and widespread methods for solving large scale linear systems of equations that arise, for example, from the discretization of partial differential equations. In this talk we introduce a new approach for optimizing the computational cost of the *full multigrid* method to achieve a given accuracy by determining the number of multigrid cycles on each level. To obtain this, a very efficient and flexible *Branch and Bound* algorithm is presented. Its implementation in the parallel finite element solver *Hierarchical Hybrid Grids* will show a significant reduction of CPU time.

Peter Thum
Towards a Physics-friendly Algebraic Multigrid Method for Systems of PDEs

Fraunhofer Institute SCAI
Schloss Birlinghoven
53754 St. Augustin
Germany
`Peter.Thum@scai.fraunhofer.de`
Tanja Clees

For complete paper, please visit
<http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/PThum.2009.pdf>

Algebraic multigrid methods (AMG) offer an efficient solution technique for many industrial applications. However, the employment of AMG for very complex applications which deal with convection, diffusion, migration and reaction terms is rather new. We focus on developing a reordering framework which is used to improve the physical-friendliness of the employed AMG approach. The presented method is predicated on point-based AMG and incorporates level-dependent alternating smoothing. Numerical results are presented for a range of electrochemical test cases with scientific and industrial relevance. The results indicate the usefulness of basing the AMG components upon the physics of the underlying system.

Guohua Zhou
**On the Accuracy of Multigrid Truncation Error Estimates on Staggered
Grids**

Department of Mathematics and Computer Science
Clarkson University
Potsdam, NY 13699-5815
zhoug@clarkson.edu
Scott R. Fulton

For complete paper, please visit
http://grandmaster.colorado.edu/~copper/2009/StudentCompetition/GZhou_2009.pdf

Multigrid methods can provide computable estimates of the truncation error for a discretized partial differential equation by comparing discretizations on grids of two different mesh sizes. [1] studied the differences between the standard formulation and a more accurate formulation for linear problems limited on non-staggered grids. This paper extends the accurate truncation error estimates to non-linear problems on staggered grids and gives applications to the Cauchy-Riemann equations, Stokes equations and Navier-Stokes equations.

REFERENCES

- [1] S. R. FULTON, *On the accuracy of multigrid truncation error estimates*, ETNA, 15(2003), pp. 29-37.

2. General Presentations

Abstracts are presented in alphabetical order by speaker's surname.

James Adler
**Nested Iteration First-Order Least Squares on Incompressible Resistive
Magnetohydrodynamics**

2075 Goss St #11
Boulder
CO 80302
`james.adler@colorado.edu`

Magnetohydrodynamics (MHD) is a single-fluid theory that describes Plasma Physics. MHD treats the plasma as one fluid of charged particles. Hence, the equations that describe the plasma form a nonlinear system that couples Navier-Stokes with Maxwell's equations. To solve this system, a nested-iteration-Newton-FOSLS-AMG approach is taken. The goal is to determine the most efficient algorithm in this context. One would like to do as much work as possible on the coarse grid including most of the linearization. Ideally, it would be good to show that at most one Newton step and a few V-cycles are all that is needed on the finest grid. This talk will develop theory that supports this argument, as well as show experiments to confirm that the algorithm can be efficient for MHD problems. Currently, a reduced 2D time-dependent formulation is studied. These equations can simulate a "large aspect-ratio" tokamak, with non-circular cross-sections. Here, the problem was reformulated in a way that is suitable for FOSLS and FOSPACK. This talk will discuss two test problems from this formulation: the Tearing Mode instability and the Island Coalescence Instability.

Travis Austin
**Automatic Construction of Sparse Preconditioners for High-Order Finite
Element Problems**

Tech-X Corporation
5621 Arapahoe Ave
Suite A
Boulder
CO 80303
`austin@txcorp.com`
Marian Brezina
University of Colorado at Boulder
Thomas Manteuffel
University of Colorado at Boulder
John Ruge
University of Colorado at Boulder

In recent years, the advantage to using high-order finite element methods and spectral element methods for the discretization of partial differential equations has become fully realized. The accuracy that can be achieved with high-order methods relative to the work required is more attractive than with first-order finite element methods. Furthermore, in climate modeling, researchers want a high-degree of precision at each time step in order to minimize the accumulation of errors due to long time

integrations, and in MHD modeling, strong magnetic field anisotropies are not accurately resolved without high-order finite element methods. As a result of this increased interest in high-order finite element methods, it is worthwhile to reconsider the development of optimized multigrid solution methods for systems derived from high-order finite elements and spectral elements.

High-order finite elements yield much denser systems of equations requiring greater memory consumption for both the matrix and for the corresponding solver infrastructure, like AMG. For large scale MHD calculations, it has been observed that memory is a limit due to the memory consumption of the high-order matrices. Thus, there is a need for a sparser approximation of the high-order finite element systems that still yields reasonable convergence. A well-known method for generating a sparser preconditioner is to use low-order finite elements to generate a sparse approximation. In this talk, we show that these sparse preconditioners require less memory and can produce better convergence behavior for 3D problems when inverted with an AMG method. We also introduce our concept for automatically constructing these preconditioners and address the computational cost and memory consumption of such an approach.

Allison H. Baker
Algebraic Multigrid for Elasticity

Lawrence Livermore National Laboratory
Box 808
L-560
Livermore
CA 94551-0808
abaker@llnl.gov
Tzanio V. Kolev
Lawrence Livermore National Laboratory
Ulrike Meier Yang
Lawrence Livermore National Laboratory

We are interested in the efficient solution of large systems of PDEs arising from elasticity applications. When solving linear systems derived from systems of PDEs with AMG, two accepted approaches are treating variables corresponding to the same unknown separately (the "unknown" approach) and treating variables corresponding to the same physical node together (the "nodal" approach). While the unknown approach is typically chosen for elasticity problems, convergence is often far from ideal.

In this talk, we investigate improving the interpolation of the rigid body modes. In particular, we propose extending the AMG interpolation operator to exactly interpolate the rotational rigid body modes by adding additional degrees of freedom (dofs) at each node. Our approach is an unknown-based approach that builds upon any existing AMG interpolation strategy and requires nodal coarsening. The approach fits easily into the AMG framework and does not require any matrix inversions. We demonstrate the effectiveness on several 2D and 3D elasticity problems.

Andrew Barker
**Multilevel Schwarz preconditioning for fluid-structure interaction with
application to blood flow**

Department of Applied Mathematics
526 UCB
University of Colorado
Boulder
CO 80309-0526
`andrew.barker@colorado.edu`
Xiao-Chuan Cai

Department of Computer Science University of Colorado, Boulder

Simulation of fluid-structure interaction is a complex problem that involves modeling different physics for the fluid and the structure and coupling them together in a stable and efficient manner, and developing scalable numerical methods for this highly nonlinear problem is a computational challenge. Here we develop and test parallel, scalable techniques in the multilevel Newton-Krylov-Schwarz family for solving the nonlinear, monolithically coupled fluid-structure interaction system on dynamic moving finite element meshes in the arbitrary Lagrangian-Eulerian framework. We develop and numerically evaluate parallel nonlinear and linear grid-sequencing techniques and linear multilevel preconditioners for the monolithically coupled problem with fully implicit time-stepping, compare these results to the corresponding one-level algorithm, and present applications of the method to the simulation of blood flow.

Hisham bin Zubair
**Multigrid preconditioning for the Indefinite Helmholtz equation on
locally refined grids**

Mekelweg 4
2628 CD
Delft
The Netherlands
`h.binzubair@tudelft.nl`

Efficient solution of the indefinite Helmholtz equation remains a challenge for multigrid based numerical solution techniques. Multigrid preconditioning techniques based on the complex shifted Helmholtz operator exist, which we exploit in this work. We deal with specific Helmholtz equations, arising out of applications from particle physics. Accuracy requirements constraint the maximum mesh size to obey $kh < 0.625$ limit. In our model Helmholtz equations, the wave number has spatial dependence, which allows us to restrict the maximum mesh size limit to certain spatially chosen areas within the domain; elsewhere we discretize with larger mesh sizes. In this talk we demonstrate multigrid preconditioning based solution of these equations. The emphasis is on the L-shaped coarsening techniques developed for such locally refined grids, which allow standard coarsening throughout the domain, in contrast with existing MLAT techniques, which (usually) only coarsen in the refined subdomain. This multigrid method is based on the simplest possible components, and the resulting convergence is good. We present the obtained results in this talk.

Markus Blatt
A h-p Algebraic Multigrid Method for Discontinuous Galerkin Methods

Interdisciplinary Center for Scientific Computing (IWR)
University Heidelberg
Im Neuenheimer Feld 368
D-69120 Heidelberg
Germany
Markus.Blatt@iwr.uni-heidelberg.de

Recently the application of discontinuous Galerkin methods to elliptic problems has gained renewed interest. Consequently, efficient solution methods for the arising discrete problem are a field of active research.

In this talk an algebraic multigrid method for discontinuous Galerkin discretizations will be discussed. The coarsening approach used is a hybrid one. For the first coarse level an auxiliary (less complex) function space is used, namely, the space of piecewise linear functions. At this level the coarsening scheme is switched and aggregation algebraic multigrid is used to construct the other coarser levels.

Matthias Bolten
Local Fourier analysis of interpolation operators for problems with certain complex stencils – preliminary results

University of Wuppertal
Department of Mathematics and Science
D-42097 Wuppertal
Germany
bolten@math.uni-wuppertal.de

Local Fourier analysis (LFA) is known to be a valuable tool for the development of geometric multigrid methods for PDEs allowing an effective tuning of multigrid components. The main concept of LFA is the idea of keeping local stencils fixed and treating the (local) problems as if they were part of the associated infinitely large constant coefficient problems. This allows for an optimal choice of smoothers as well as restriction operators.

The infinitely large problem with constant coefficients corresponds to a (l -level) Toeplitz operator. The Toeplitz operators are completely described by their generating symbols, i.e. (l -variate) 2π -periodic functions. For second order elliptic problems the generating symbol has a unique zero of order two at the origin. The multigrid methods for Toeplitz matrices and circulant matrices that have been developed in the last years work well for these problems, and they do not depend on the location of the zero. In fact, the zero of the generating symbol just influences the choice of the interpolation operator in the multigrid method.

Using these developments we are able to provide local interpolation operators for matrices with non-constant stencils with complex entries, consider e.g. the stencil

$$\frac{1}{h^2} \begin{bmatrix} & -e^{2\pi i\varphi_{x,y+\mu}} & \\ -e^{-2\pi i\varphi_{x-\mu,y}} & 4 & -e^{2\pi i\varphi_{x+\mu,y}} \\ & -e^{-2\pi i\varphi_{x,y-\mu}} & \end{bmatrix}$$

where the $\varphi_{x-\mu,y}, \varphi_{x+\mu,y}, \varphi_{x,y-\mu}, \varphi_{x,y+\mu} \in [0, 1]$ are some random parameters. This can be considered as a non-shifted variant of the 2d Gauge Laplace matrix that arises in a simplified model of lattice quantum chromodynamics. The idea is similar as in LFA, namely take the local stencil as a constant stencil in an infinitely large system. Using results found in previous works on multigrid methods for matrices being Toeplitz or circulant we are able to provide a local definition of the interpolation, yielding in interpolation operators for these matrices. The method has been implemented for the two-grid case and the results are as good as expected. The idea can be used in a multilevel setting, easily.

In this talk we will give a short overview over the used previous work, introduce the concept in larger detail and present some numerical results for the two-grid case.

S. D. Buckeridge
Parallel Multigrid in Numerical Weather Prediction

University of Bath
 Department of Mathematical Sciences
 Bath
 BA27AY
 UK
 sdb22@bath.ac.uk
 R. Scheichl
Supervisor

Fast elliptic solvers are crucial to the forecasting and data assimilation tools used at the UK Met Office - one of the world's leading providers of weather-related services. Some of the core equations used in the numerical solution of weather and climate models are the Quasi-Geostrophic Dynamical equations, in particular the Quasi-Geostrophic Omega equation

$$N^2(r)\nabla^2 u - f_0^2 \frac{1}{r^2} \frac{\partial^2 u}{\partial r^2} = f.$$

The Met Office uses structured spherical polar grids which has the drawback of creating strong anisotropies near the poles where the grid lines converge. Moreover, the grid spacings in the radial direction are much smaller since the thickness of the atmosphere is small compared to the circumference of the Earth's surface. Additionally, the grid is graded in this direction, with smaller grid spacings near the surface of the Earth to obtain a better resolution in the regions of most interest, thus creating a strong anisotropy also in the radial direction. Multigrid methods are known for their efficiency and robustness for isotropic elliptic problems, and remedies exist for anisotropic problems such as semi-coarsening and line smoothing. However, since the strength of anisotropy varies between the equator and the poles, the existing methods must be adapted further, introducing a non-uniform coarsening strategy, where the grid is coarsened only in regions that are sufficiently isotropic.

The success of non-uniform coarsening strategies has been demonstrated with Algebraic Multigrid (AMG) methods. Without the large setup costs required by these methods, however, we aim to surpass them with the geometric approach outlined above. Results will be given for both sequential and parallel solves.

Mathieu Chanaud
**A dedicated method for solving large linear systems arising from finite
element discretizations of high-frequency Maxwell problems.**

INRIA Bordeaux Sud-Ouest and LaBRI
Université de Bordeaux
Talence
France

chanaud@labri.fr

David Goudin

CEA/CESTA, BP 2, 33114 Le Barp, France

Jean-Jacques Pesqué

CEA/CESTA, BP 2, 33114 Le Barp, France

Jean Roman

INRIA Bordeaux Sud-Ouest and LaBRI, Université de Bordeaux, Talence, France

Solving Maxwell equations using multigrid method is a well known problem, and many approaches and solvers have been explored. In this talk, we describe a specific methodology for solving large linear systems arising from Maxwell equations discretized with first-order Nédélec elements. Using this methodology, a new parallel solver is developed. This solver combines the PaStiX parallel direct solver and full multigrid cycles. The goal of this proposed method is to compute the solution for problems defined on fine irregular meshes with minimal overhead costs when compared to the cost of applying a classical direct solver on the coarse mesh.

The PaStiX direct solver can handle linear systems with up to 100 million unknowns, but this size is limited by the computer memory, so that finer problem resolutions that often occur in practice cannot be handled by this direct solver. The aim of the new method is to provide a way to solve problems with up to 1 billion unknowns, given an input mesh with up to 100 million unknowns. The input mesh is used as the initial coarsest grid. Fine meshes, where only smoothing sweeps are performed, will be generated automatically. Such large problem size can be solved because the matrix is assembled only on the coarsest mesh, while on the finer meshes, multigrid cycles are performed matrix-free.

The problem is defined on a 3D domain, which might be a subdomain extracted from a physical domain Ω using a domain decomposition technique. There are two types of boundaries, giving rise to two types of boundary conditions, an impedance condition on the inner, material interface and a transmission condition on the outer interface.

In order to solve the differential equations, we use a tetrahedral Nédélec finite element discretization, leading to an irregular 3D mesh. These elements are edge-based, with the unknowns defined on the edges. These edge values represent the electric field circulation along an oriented edge Γ_i with unit tangential vector $\vec{\tau}_i$, and they are computed in each element with the formula

$$\alpha_i = \int_{\Gamma_i} \vec{E} \cdot \vec{\tau}_i d\Gamma_i.$$

In each element, electric field interpolation \hat{E} is achieved using the natural linear vectorial basis

function p_i

$$\begin{aligned} \int_{\Gamma_j} \vec{p}_i \cdot \vec{\tau}_j d\Gamma_j &= 0 \text{ for } \Gamma_j \neq i \\ &= 1 \text{ for } \Gamma_j = i, \\ \hat{E}(x, y, z) &= \sum_{i=1}^6 \vec{p}_i(x, y, z) \cdot \alpha_i. \end{aligned}$$

With these components, we can show that mesh refinement is numerically simple, and two regular refinement methods can be used. The first refinement method uses the gravity center as a new vertex, creating 4 new tetrahedrons. The new edge values, computed using Nédélec basis functions, are simply weighted sums of the coarser edge values, which are not modified by this refinement procedure. The second method splits each tetrahedron edge. As a consequence, the coarsening phase for this second method is more complicated and may require more memory storage to keep vectors on all mesh levels. Furthermore, in this second method, for each coarse tetrahedron, 8 new ones are created, which means that problem size grows fast on each refinement level. Because of this growth, we use the first refinement method.

The first step of our method is to solve the problem on the coarsest input mesh using the PaStiX parallel direct solver. On this coarsest level, the system is formed by assembling the global matrix using the elementary matrices. This complex but symmetric matrix is then factorized and solved. Next, the computed solution is prolonged onto a finer mesh using the technique described above. On this finer mesh, a multigrid V-cycle is used to damp out the interpolation error and capture fine-scale solution features.

Choosing the smoother in the V-cycle is crucial. Such method must smooth the error in a few sweeps, and must also be efficient and parallelizable. Moreover, in order to be matrix-free, the smoother must be realized using only matrix-vector products. Currently, we plan to use Gauss-Jacobi type smoothers.

Since the input is the coarsest mesh, the depth of the multigrid cycle is determined by this original mesh level. In the V-cycle, once the coarsest refinement level is reached, a direct solve is performed using the factorized matrix already computed. As a consequence, solving the coarsest level system requires only a forward-backward solve phase. Naturally, reaching this coarse level is not always necessary if the desired accuracy has already been attained.

The main advantage of this method is that the coarsest level matrix only needs to be assembled, and its factorization is done only once (since we use a regular refinement procedure, the coarsest level matrix is unchanged). Factorization performed by the first PaStiX invocation can be viewed as a preprocessing stage, piloting the parallel aspects of the solution phase by distributing the matrix and elements over processors. All solves after the parallel factorization are performed with a complexity equivalent to a forward-backward solve, which is highly dependent on the renumbering of the unknowns performed before the factorization.

Another advantage of this method is related to the mesh. In classical Maxwell problems, the mesh must meet two requirements. It must physically describe the problem and it must satisfy a wavelength restriction that ensures that elements are small enough to capture the applied wavelength. In this approach, the input mesh must fulfill the first requirement, while the second one only applies to the fine mesh and is progressively ensured by successive refinements. From this specificity arises one main advantage of this method. Studying a given object at increasing high frequencies requires a unique initial mesh describing problem geometry while wavelength requirements are met by using deeper V-cycles.

Cy P Chan
A Dynamic Programming Approach to Autotuning Multigrid

44 Winter St #1
Cambridge
MA 02141

cychan@csail.mit.edu

Jason Ansel

*Massachusetts Institute of Technology, Department of Electrical Engineering and Computer
Science, Computer Science and Artificial Intelligence Laboratory*

Yee Lok Wong

Massachusetts Institute of Technology, Department of Mathematics

Saman Amarasinghe

*Massachusetts Institute of Technology, Department of Electrical Engineering and Computer
Science, Computer Science and Artificial Intelligence Laboratory*

Alan Edelman

*Massachusetts Institute of Technology, Department of Mathematics, Computer Science and
Artificial Intelligence Laboratory*

Algorithmic choice is essential in any problem domain to realizing optimal computational performance. Multigrid is a prime example: not only is it possible to make choices at the highest level, but a program can switch techniques as the problem is recursively attacked on coarser grid levels to take advantage of algorithms with different scaling behaviors. Additionally, users with different convergence criteria must experiment with parameters to yield a tuned algorithm that meets their accuracy requirements. Even after a tuned algorithm has been found, users often have to start all over when migrating from one machine to another.

We present a programming language and autotuning system that addresses all of these issues in a near-optimal and efficient fashion. The freedom of independently tuning both the algorithm and the number of iterations at each recursion level results in an exponential search space of hybrid algorithms that have different accuracies and performances. To search this space efficiently, our autotuner utilizes a novel dynamic programming method to build near-optimal hybrid algorithms from the bottom up. The result is a tuned hybrid algorithm that invests minimal targeted computational power to yield the accuracy required by the user.

The techniques we describe allow the user to automatically generate tuned "W-cycles" of different shapes targeted to the user's specific combination of problem, hardware, and accuracy requirements. These cycle shapes dictate the order in which grid coarsening and grid refinement are interleaved with both iterative methods, such as Jacobi or Successive Over-Relaxation, as well as direct methods, which tend to have superior performance for small problem sizes. The need to make choices between all of these methods brings the issue of variable accuracy to the forefront. Not only must the autotuning framework compare different possible W-cycle shapes against each other, but it also needs the ability to compare W-cycles against both direct and (non-multigrid) iterative methods. We solve this problem by using accuracy as a measure for the effectiveness of tuned W-cycle shapes and making comparisons over all algorithmic types based on this common yardstick. In our results, we find that the flexibility to trade performance versus accuracy at all levels of recursive computation enables us to achieve outstanding performance on a variety of platforms compared to algorithmically static implementations of multigrid.

Our implementation uses PetaBricks, an implicitly parallel programming language where algorithmic choices are exposed in the language. The PetaBricks compiler uses these choices to analyze, autotune, and verify the PetaBricks program. These language features, most notably the autotuner, were key

in enabling our implementation to be clear, correct, and fast.

Michael A Clark
Adaptive multigrid algorithm for the lattice QCD Wilson-Dirac matrix

Center for Computational Sciences
Boston University
3 Cummington St
Boston
MA 02215
mikec@bu.edu
James Brannick
Penn State
Richard Brower
Boston University
Tom Manteuffel
University of Colorado at Boulder
Steve McCormick
University of Colorado at Boulder
James Osborn
Argonne National Laboratory
Claudio Rebbi
Boston University

The Wilson-Dirac matrix is the discretized Dirac operator that describes fermions in quantum field theory. The parameter of interest in this matrix is the fermion mass parameter; the matrix becomes singular as this parameter is taken to zero. Solving such systems of linear equations at near zero fermion mass is the prevalent computational cost in lattice QCD calculations.

Although the matrix has a regular geometric structure, the random nature of the coefficients of the matrix results in the failure of naive multigrid approaches. A geometric adaptive multigrid algorithm based on adaptive smooth aggregation (α SA) has been developed specifically for this matrix. Special consideration must be paid to the smoother used since the matrix is not Hermitian positive definite. The resulting algorithm exhibits very little dependence on the condition number of the matrix.

Marco, M Donatelli
An algebraic generalization of the local Fourier analysis

Universita' dell'Insubria - sede di Como
Dipartimento di Fisica e Matematica
via Valleggio
11 - 22100 Como (ITALY)
marco.donatelli@uninsubria.it

The Local Fourier analysis (LFA) is a classic tool to prove convergence results for multigrid methods (MGMs). In particular, we are interested in optimality that is a convergence speed independent of the size of the involved matrices. For elliptic partial differential equations (PDEs), a well known optimality result requires that the sum of the orders of the grid transfer operators is not lower than the order of the PDE to be solved. Analogously, when dealing with MGMs for Toeplitz matrices in

the literature an optimality condition on the position and on the order of the zeros of the symbols of the grid transfer operators has been found. In this work we show that in the case of elliptic PDEs with constant coefficients, the two different approaches lead to an equivalent condition. We argue that the analysis for Toeplitz matrices is an algebraic generalization of the LFA, which allows to deal not only with differential problems, but also, for instance, with integral problems. We give also a class of grid transfer operators related to the B-spline refinement equation and we study their geometric properties. This analysis suggests further links between wavelets and multigrid methods. A numerical experimentation confirms the correctness of the proposed analysis.

Craig C. Douglas
**AMGLab: Algebraic Multigrid in Matlab to Production Computing on
Parallel Supercomputers**

University of Wyoming
Department of Mathematics
1000 E University Ave
Dept 3036
Laramie
WY 82071
craig.c.douglas@gmail.com
Craig C. Douglas
University of Wyoming

AMGLab is an open source, scientific community oriented project to allow nonexperts to experiment with AMG algorithms and then be able to choose a robust parallel AMG solver. Industrial class problems are the target. However, student level problems, suitable for classes, are also a target. The project has been centered at the University of Kentucky, University of Graz, and the University of Linz. The speaker has recently moved to the University of Wyoming.

In this talk, we will discuss what is implemented in Matlab, what exists for production level parallel codes, and how others can become part of the project by contributing to the open source project.

Jose Garcia
Comparing Multigrid Solver Performance on Many-Core Accelerators.

PO Box 3000
Boulder
CO
80307
jgarcia@ucar.edu
Jose Garcia
National Center for Atmospheric Research (NCAR). University of Colorado
Rory Kelly
National Center for Atmospheric Research (NCAR).

Multigrid solvers are considered to be the most efficient methods to solve many classes of numerical problems. However, the true performance of these methods depends on their specific implementations for different computer architectures. In recent years the continuing increase in single threaded micro-processor performance has stalled as designs have become bound by fundamental limitations with the

underlying device physics. These limitations have ushered in a new era of multi-core, many-core, and heterogeneous core processors, aimed at increasing speed through explicit parallelism. For Multigrid algorithms to benefit from this recent trend they must be able to efficiently exploit the parallelism offered by the underlying hardware. We explore the complexities, performance and limitations that Multigrid methods may face when implemented for several current generation accelerators, including the Cell Broadband Engine and NVIDIA Graphics Processing Units.

Francisco Gaspar
Design of geometric multigrid methods on semi-structured grids

Department of Applied Mathematics
University of Zaragoza
c/Maria de Luna 3
50018 Zaragoza
Spain

`fjgaspar@unizar.es`

Jose Luis Gracia
University of Zaragoza
Francisco Lisbona
University of Zaragoza
Carmen Rodrigo
University of Zaragoza

We are interested in the design of efficient geometric multigrid methods on hierarchical triangular grids for problems in two dimensions. Assuming that the coarsest grid is rough enough in order to fit the geometry of the domain, a hierarchy of globally unstructured grids is generated. This kind of meshes are suitable for use with geometric multigrid. To discretize problems with constant coefficients on these type of meshes, explicit assembly of the global stiffness matrix for the finite element method is not necessary and this can be implemented using stencils. As the stencil for each coarsest triangle is the same for all unknowns that are interior to it, one stencil suffices to represent the discrete operator reducing drastically the memory required. Fourier analysis is a well-known useful tool in multigrid for the prediction of two-grid convergence rates. With the help of the Fourier Analysis on triangular grids, we design efficient geometric multigrid methods for different problems on hierarchical triangular grids.

Tobias Gradl
High Performance Adaptive Mesh Refinement

Universitaet Erlangen-Nuernberg
Lehrstuhl fuer Informatik 10
Cauerstr 6
91058 Erlangen
Germany

`tobias.gradl@cs.fau.de`

Ulrich Ruede

Department of Computer Science 10, University of Erlangen-Nuremberg

Getting multigrid software to solve for as many unknowns per second as possible on current high performance computers means designing its data access patterns to be as regular as possible. *Regu-*

larly structured data allows for optimally exploiting the processor capabilities (e. g. vector processing) and the interconnection network (e. g. by grouping communication operations). Just as important as solving for many unknowns in short time is reducing the overall number of unknowns. *Adaptive mesh refinement (AMR)* can significantly reduce the problem size in many cases. However, this technique is known to impede the use of regular data structures, and, thus, to drastically restrict the possible performance of the solver. In our talk we show how AMR can be implemented without destroying the regularity of the data structures, maintaining a fast execution. The framework for the AMR implementation is provided by *Hierarchical Hybrid Grids (HHG)*, a multigrid finite element solver designed for the use on massively parallel computers with several thousand processors. The software uses semi-structured meshes in order to achieve high performance on such computers. In the talk both red-green refinement and refinement with hanging nodes are covered, because both methods have individual advantages, and the decision for one of the methods—or for a combination of both—has to be made depending on the problem characteristics. We develop the mathematical foundation for using AMR with the *correction scheme*, and we show its feasibility in practical examples.

Jeffrey J. Heys
**Algebraic Multigrid Solvers for Weighted Least-Squares Finite Elements
with Application to Particle Imaging Velocimetry Analysis**

Chemical and Biological Engineering
Montana State University
PO Box 173920
Bozeman
MT 59717-3920
jeff.heys@gmail.com
Michele Milano
Arizona State University
Marek Belohlavek
Mayo Clinic Scottsdale

The combination of ultrasound and microbubbles injected into the blood makes it possible to non-invasively obtain 2-dimensional velocity data inside the left ventricle of the heart. A long term goal is to translate this velocity data into information about energy loss, pressure gradients, and overall health of the heart, but this goal requires a full 3-dimensional velocity field. The question that we seek to answer is whether the 2-dimensional velocity data can be incorporated into a full 3-dimensional computational fluid dynamics simulation in an appropriate and computationally practical way. For addressing this problem, we examine the potential of least-squares finite element methods (LSFEM) because of their flexibility in the enforcement of various boundary conditions and their natural compatibility with multigrid solvers. By weighting the boundary conditions in a manner that properly reflects the accuracy with which the boundary values are known, we develop the weighted LSFEM. The potential of weighted LSFEM is explored for two different test problems: the first uses randomly generated Gaussian noise to create artificial ‘experimental’ data in a controlled manner, and the second uses experimental particle imaging velocimetry data. In both test problems, weighted LSFEM produces accurate results even for cases where there is significant noise in the experimental data and provides excellent computational scalability when used with a parallel algebraic multigrid solver.

Lior Horesh
Refinement criteria for OcTree discretization of Maxwell equations

Department of Mathematics and Computer Science
Emory University
400 Dowman Dr
Atlanta
GA
30322
horesh@mathcs.emory.edu
Eldad Haber

In this study an OcTree discretization for Maxwell's equations is proposed. This discretization is based on an explicit construction of the adjoint of the curl operator. As opposed to the $\mathcal{O}(1)$ local accuracy offered by the transposed curl operator, this non-symmetric construction offers local accuracy of $\mathcal{O}(h)$, which is more favorable. We explore several adaptive refinement criteria, and discuss their effectiveness.

Raya Horesh
A Multi-grid framework for volume preserving constrained image registration on OcTree structure

Department of Mathematics and Computer Science
Emory University
400 Dowman Dr
Atlanta
GA 30322
USA
rshindm@emory.edu
Eldad Haber
Jan Modersitzki

Image registration is one of the fundamental missions of image processing. It can be simply considered as a process of aligning/matching two or more images having similar contents in some sense. For example, the images could have been captured at different times, from different viewpoints and/or using different types of sensors. Since the problem of image registration is ill-posed, one may wish to add additional information. In this study, the deformation is controlled in terms of the determinant of the Jacobian of the transformation. This approach guarantees regularity of the grid and prevent folding effects.

In order to keep the computational time reasonable, it is desirable to reduce the amount of data which need to be processed. This can naively be obtained by down-sampling. However, this approach may results in the loss of possibly important image features. OcTree provides a straightforward approach for such a context driven data sparsification. The OcTree data presentation uses few voxels for representing regions with small variability and many voxels for regions of high variability.

Discretization of the volume preserving constrained registration problem yields a KKT system. This large-scale system is indefinite and ill-conditioned; therefore, we propose a multi-grid framework for effective preconditioning of this system. We present mesh-size-independent results that demonstrate the optimality of the proposed method.

Gary W. Howell
BLAS-3 Sparse $UB_{K+1}V$ Decomposition

512 Farmington Woods Dr
Cary
NC 27511
gary_howell@ncsu.edu

A matrix A can be reduced to upper triangular banded form by BLAS-3 block Householder transformations. Deferring matrix updates, the algorithm accesses A only to extract blocks and to perform multiplications AX , AY^T , where if X and Y have K columns than the bandwidth of upper triangular B_{K+1} is $K + 1$. When A is sparse, block Householder eliminations provide a BLAS-3 method to construct a $UB_{K+1}V$ approximation, with U and V orthogonal, U ($m \times l$), V ($n \times l$), B_{K+1} $l \times l$, with the size of l constrained by available RAM.

The decomposition is stable, is efficient in terms of cache utilization, and should scale well in distributed parallel computation.

The $UB_{K+1}V$ approximate (truncated) decomposition can be used to provide some matrix singular values, to solve linear systems and least squares problems, and to provide an approximate inverse preconditioner. Multi-grid applications may be parallel iterations with the full matrix, as a preconditioner, or in solution of a coarsened problem. Each of these applications is discussed.

A primary advantage of the block reduction to a banded upper triangular form is that the underlying sparse matrix is accessed only for multiplications by blocks of matrices (sparse matrix dense matrix multiplications). Serial SPMD multiplications run at a significant fraction of peak speed on cache based processors, and should also run well in parallel. Stability of block Householder transformations aids in scalability.

Numerically, the truncated $UB_{K+1}V$ decomposition is seen to be particularly efficient in approximating low rank matrices or low rank matrices added to a matrix with a known factorization.

Some unresolved questions are how to best prepermute A , how to best pad A (thick start), and considered as a Krylov method the best restart strategies (thick restart?, repermuation of A ?).

The remainder of the abstract discusses why multiplication of AX is likely to be faster than computing Ax , assuming A sparse, x a dense vector X a dense matrix with relatively few columns. Assume A is too large to fit in cache memory. A is typically stored so that it can be pulled in a stream from RAM. If x fits in cache, then as each element of A appears in the CPU it can be matched by the appropriate element of x . If indexing operations do not take too long the main cost is then the fetch of A from RAM. Since the fetch of A is streamed, elements of A arrive more or less at the peak speed of the data bus.

On Intel Xeons, sparse Ax can attain up to about ten per cent of the advertised peak flop rate. When x becomes too large to fit in cache, and if x is accessed randomly, then the computation is dominated by cache misses and slows dramatically. In some experiments with Intel Xeons, Ax computed at less than one per cent of peak processor speed. When multiplying by X with K columns as opposed to x , each access of an element of A allows K multiplications. Storage of X should be arranged so that when a given element of A is accessed, the next required elements of X are accessed. In Fortran, X for AX is stored as X^T so that each row of X is sequentially stored as a column of X^T .

In experiments summarized here, we also blocked A (column blocks) to further minimize cache misses. The blocked SPMD AX can execute at a flop rate several orders of magnitude faster than

non-blocked Ax . The marked speedup of BLAS 3 over BLAS 2 motivates the algorithm development described in the presentation.

Kirk E. Jordan

Title: Petascale Barrier Surpassed But How Will We Solve Real Problems on Petascale Systems and/or Exascale Systems?

IBM T J Watson Research Center

<http://www-3.ibm.com/software/info/university/people/kjordan.html>

kjordan@us.ibm.com

High performance computing (hpc) is a tool frequently used to understand complex problems in numerous areas such as aerospace, biology, climate modeling and energy. Scientists and engineers working on problems in these and other areas demand ever increasing compute power for their problems. In order to satisfy the demand for increase performance to achieve breakthrough science and engineering, we turn to parallelism through large systems with multi-core chips. For these systems to be useful massive parallelism at the chip level is not sufficient. I will describe some of the challenges that will need to be considered in designing Petascale and eventually Exascale systems. However, the hardware development is not as hard as designing algorithms that will exploit these systems. Through the combination of hpc hardware coupled with novel algorithmic approaches, such as multi-grid methods, some efforts toward breakthroughs in science and engineering are described. While progress is being made, there remain many challenges for the computational science community to apply ultra-scale, multi-core systems to “Big” science problems with impact on society. In conclusion, some discussion not only on the most obvious way to use ultra-scale, multi-core hpc systems will be given but also some thoughts on how one might use such systems to tackle previously intractable problems.

Ralf Juengling

Multigrid vs. Conjugate Gradient for Piece-Wise Smooth Surface Reconstruction, a Case Study

Portland State University

Department of Computer Science

P O Box 751

Portland

OR 97207-0751

juenglin@cs.pdx.edu

An interesting problem from the field of computer vision is reconstruction of piece-wise smooth functions from noisy and incomplete data. The problem is more difficult than reconstructing smooth functions and algorithms typically employ an iterative relaxation method as a component.

We studied a particular algorithm, published twenty years ago by Leclerc [1], which reconstructs two-dimensional, piece-wise polynomial functions from noisy data. An objective function is derived from the premise that a description of the data in terms of a piece-wise polynomial signal plus noise from some known stochastic process is more likely to be correct the shorter it is. The objective, f , is not continuous and thus difficult to optimize.

Leclerc devises a continuation method by which a sequence of tentative solutions is sought, solutions corresponding to a sequence of smooth objective functions converging to f . For each step of the

continuation method Leclerc uses the Jacobi method to relax the most recent tentative solution on the linear system that results from setting the gradient of the next objective in the sequence to zero.

In our experiments it turned out that Leclerc's algorithm finds satisfying solutions only for small problem sizes and performs poorly with realistic sizes. In several variations of Leclerc's continuation scheme we substitute other methods for the Jacobi relaxation to remedy the poor performance. In particular, we substitute a Conjugate Gradient method, an inexact Newton method, and a multigrid Gauss-Seidel relaxation method.

In this presentation we briefly review Leclerc's work and give a picture of the difficulties we encountered with larger problem sizes. We sketch our variations of Leclerc's continuation algorithm, present results for each method and discuss their merits and shortcomings in terms of accuracy, computational efficiency, and amount of work required for their implementation.

REFERENCES

- [1] Y. Leclerc: "Constructing simple stable descriptions for image partitioning", IJCV (3), pp. 73-102, 1989.

Christian Ketelsen
Multigrid Solvers for Quantum Electrodynamics

Department of Applied Mathematics
526 UCB
University of Colorado
Boulder
CO 80309-0526
`ketelsen@colorado.edu`

Tom Manteuffel
University of Colorado at Boulder
Steve McCormick
University of Colorado at Boulder

The Schwinger Model of quantum electrodynamics (QED) describes the interaction between electrons and photons. Large scale numerical simulations of the theory require repeated solution of the two-dimensional Dirac equation, a system of two first order partial differential equations coupled to a background U(1) gauge field. Traditional discretizations of this system are sparse and highly structured, but contain random complex entries introduced by the background field. For even mildly disordered gauge fields the near kernel components of the system are highly oscillatory, rendering standard multilevel methods ineffective.

We consider an alternate formulation of the governing equations that is more amenable to multigrid solvers. The alternate model is obtained by a similarity transformation of the continuum operator which essentially eliminates the gauge field. The resulting formulation resembles uncoupled diffusion-like problems with variable diffusion coefficients. The form of the transformed system is ideal as adaptive multilevel methods have proven effective in the past at solving such problems. Next, we discretize the transformed system using least-squares finite elements. Finally, adaptive smoothed aggregation multigrid is used to solve the resulting linear system. We present numerical results and discuss implications of the transformed formulation for the physical theory.

Andrew Knyazev
Multigrid Eigensolvers for Image Segmentation

Department of Mathematical and Statistical Sciences
University of Colorado Denver
P O Box 173364
Campus Box 170
Denver
CO 80217-3364
`Andrew.Knyazev@ucdenver.edu`

Known spectral methods for graph bipartition and graph-based image segmentation require computation of Fiedler vectors, i.e., numerical solution of eigenvalue problems with a graph Laplacian. The ultimate goal is to find a method with a linear complexity, i.e. a method with computational costs that scale linearly with the problem size, e.g., with the number of pixels in the image for the image segmentation problem. Multigrid approaches seem natural for image segmentation, where different image resolution scales are easily available.

We numerically analyze multigrid-based eigensolvers, e.g., [1-2], for computation of the Fiedler vectors. Multigrid can be used for multi-resolution segmentation as well as for preconditioning. We test both such approaches. We find that the multiresolution segmentation can be tricky as the low-resolution image bisection may be qualitatively inaccurate; and we explain the mathematical reason for such a behavior. A direct bisection of the highest-resolution image may thus produce better quality segmentations compared to the multiresolution segmentation.

Our tests demonstrate that the multigrid preconditioning gives the ideal linear complexity and produces high quality image segmentation applied directly to the highest-resolution image. We describe our PETSc-BLOPEX, see [3], driver for computing the Fiedler vector with Hypre preconditioning, which can be used for segmentation of practical-size megapixel images on parallel computers. E.g., it computes the Fiedler vector for 24 megapixel images in seconds on our BlueGene/L 1024 CPU box using the algebraic multigrid. Further speed-up can be obtained by employing geometric multigrid and a lower precision arithmetic.

REFERENCES

- [1] A.V. Knyazev, Toward the Optimal Preconditioned Eigensolver: Locally Optimal Block Preconditioned Conjugate Gradient Method. *SIAM Journal on Scientific Computing* 23 (2001), 517-541.
- [2] A.V. Knyazev and K. Neymeyr, Efficient solution of symmetric eigenvalue problems using multigrid preconditioners in the locally optimal block conjugate gradient method. *ETNA*, 15 (2003), 38-55.
- [3] A. V. Knyazev, I. Lashuk, M. E. Argentati, and E. Ovchinnikov, Block Locally Optimal Preconditioned Eigenvalue Solvers (BLOPEX) in hypre and PETSc. *SIAM Journal on Scientific Computing* 25 (2007), 2224-2239.

Harald Koestler
**A Multigrid Prolongation Method Based on Local Minimization of the
Constant in the Strengthened Cauchy-Schwarz Inequality**

University of Erlangen-Nuremberg
Lehrstuhl für Informatik 10
Cauerstrasse 6
D-91056 Erlangen
Germany
`harald.koestler@informatik.uni-erlangen.de`
Christoph Pflaum
University of Erlangen-Nuremberg

We discuss ways to construct prolongation operators for multigrid methods both for geometric and pure algebraic settings. Our theoretical derivation of the prolongation operator is based on finite element analysis, in detail we try to minimize the constant in the strengthened Cauchy-Schwarz inequality locally. Therefore, we first write the fine grid space as a direct sum of the coarse grid space and a suitable complementary space. The local coarse subspaces are chosen such that they can locally represent the kernel of the operator. In order to solve the arising local minimization problems and to determine the prolongation weights only a few local operations are required. We show first numerical results in 2D and outline applications of our method like linear elasticity.

Yiannis Koutis
The Combinatorial Multigrid Solver

Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh
PA 15213
`jkoutis@cs.cmu.edu`
Gary Miller
Carnegie Mellon University

The study of systems on graph Laplacians -the discrete analog of the inhomogeneous Poisson equation- has been a topic of interest for both the scientific computing (SC) and the theoretical computer science (TCS) communities. While the SC community has developed multigrid methods for this problems since the 70s, the TCS community has started studying the problem much more recently, in the middle 90s, through the introduction of Combinatorial Preconditioners, by P. Vaidya. The intense study of Combinatorial Preconditioners has recently lead to algorithms with strictly provable asymptotic properties: the Spielman-Teng solver, with a complexity of $O(m \text{ polylog}(m))$ for general Laplacians with m non-zeros, and our $O(n)$ solver for planar Laplacians.

In this talk we present the Combinatorial Multigrid Solver (CMG). The CMG solver is a multigrid solver derived through the construction of combinatorial preconditioners that are based on the combinatorial geometry of the underlying graph. In that way the CMG solver combines the "black-box/theoretical guarantees" quality of the combinatorial preconditioners approach, with the speed and the parallelism potential of multigrid. To validate our claim, we present experiments with very large 3D Laplacians derived from medical scans.

Milan Kucharik
**Application of JFNK to Kinetic Energy Conservative Remapping of the
 Nodal Velocity in ALE**

T-05 MS B284
 Los Alamos National Laboratory
 Los Alamos
 NM 87545 USA
 kucharik@lanl.gov
 Markus Berndt
 Mikhail Shashkov

Arbitrary Lagrangian-Eulerian (ALE) methods appear to be reasonable compromise between Lagrangian and Eulerian approaches, allowing to solve large variety of fluid problems. The standard ALE algorithm uses a Lagrangian solver to update fluid quantities and the computational mesh in the next time step, which can eventually tangle the mesh. To avoid such problems, mesh regularization (untangling or smoothing) is applied in the case of low mesh quality, followed by a remapping step that interpolates all fluid quantities from the Lagrangian to the smoothed mesh.

Here, we focus on the last part of the ALE algorithm – remapping – in the case of a staggered discretization, where scalar quantities (density, pressure, specific internal energy) are defined inside mesh cells, and vector quantities (positions, velocities) are defined on mesh nodes. A staggered discretization is used in most current ALE codes. Generally, remapped nodal kinetic energy is not equal to nodal kinetic energy obtained from remapped velocities, and this discrepancy leads to energy conservation violation and consequently to wrong shock speeds. The kinetic energy discrepancy is usually treated by distributing it to the internal energy of adjacent cells [2].

An alternative 1D approach introduced in [1] is based on the construction of high-order interpolated velocities $u_{n+1/2}^*$ used for momentum fluxes (and thus for the velocity update)

$$\tilde{m}_n \tilde{u}_n = m_n u_n + F_{n,n+1/2}^m u_{n+1/2}^* - F_{n,n-1/2}^m u_{n-1/2}^* \quad (1)$$

such that the kinetic energy discrepancy is automatically eliminated. Here, $F_{n,n\pm 1/2}^m$ stands for mass fluxes around node n , which are also used for the nodal mass update in a similar flux form

$$\tilde{m}_n = m_n + F_{n,n+1/2}^m - F_{n,n-1/2}^m. \quad (2)$$

This method expresses the new kinetic energy \tilde{K} in each node in a flux form

$$\tilde{K}_n = K_n + F_{n,n+1/2}^K - F_{n,n-1/2}^K, \quad (3)$$

where fluxes F^K can be computed from known quantities and unknown inter-nodal flux velocities $u_{n\pm 1/2}^*$. To achieve energy conservation, there must be an unique flux between every two nodes:

$$F_{n,n+1/2}^K = F_{n+1,n+1/2}^K \equiv F_{n+1/2}^K, \quad \forall n. \quad (4)$$

After substituting for all F^K , we obtain a system of coupled quadratic equations for all u^* ,

$$u_{n+1/2}^* = \frac{u_{n+1} \bar{u}_{n+1} - u_n \bar{u}_n - (u_{n+1}^2 - u_n^2) / 2}{\bar{u}_{n+1} - \bar{u}_n}, \quad (5)$$

where $\bar{u}_n = (u_n + \tilde{u}_n) / 2$, and the new velocities \tilde{u}_n are defined in (1). This system can either be solved with a fixed point iteration, or with a Newton solver.

In this presentation, we demonstrate that the system (5) does not always have a solution, and analyze this situation. Two alternative approaches will be introduced.

In the first approach, the flux equality (4) is not enforced strictly, but its discrepancy is minimized in a least squares sense

$$D^F = \sum_{\forall(n+1/2)} (F_{n+1,n+1/2}^K - F_{n,n+1/2}^K)^2 \quad (6)$$

by differentiating D^F with respect to all u^* and solving a system $\partial D^F / \partial u_{n+1/2}^* = 0$ for all $n + 1/2$. Thus, the kinetic energy discrepancy is minimized as much as possible, but is not guaranteed to equal zero.

In the second approach, the kinetic energy discrepancy is directly minimized. Unfortunately, zero kinetic energy discrepancy is generally satisfied by infinitely many solutions. Because flux velocities $u_{n+1/2}^*$ represent values interpolated from adjacent nodal velocities u_n and u_{n+1} , this velocity should remain bounded by these neighbor values and no under/overshoots should appear. Thus, additional terms are added to the kinetic energy discrepancy formula, which enforce these bounds during minimization

$$D^K = (\tilde{K} - K)^2 + \varepsilon K \sum_{\forall(n+1/2)} \left(\frac{1}{2} m_{n+1/2} \left(u_{n+1/2}^* - \frac{u_n + u_{n+1}}{2} \right)^2 \right). \quad (7)$$

Similarly to our first approach, the minimization is done by solving a system $\partial D^K / \partial u_{n+1/2}^* = 0$ for all $n + 1/2$. This approach finds a zero discrepancy solution that is generally different from the solution of system (4) (if it exists). In both approaches (6) and (7), the Jacobian-free Newton-Krylov solver [3] with a possible preconditioning is used.

We will also mention a generalization of these approaches to 2D logically orthogonal meshes (including corner coupling), and demonstrate their behavior in the context of an ALE hydro code for a particular fluid flow problem.

Acknowledgment

This work was performed under the auspices of the US Department of Energy by Los Alamos National Laboratory under Contract DE-AC52-06NA25396. The authors acknowledge the partial support of the DOE Advance Simulation and Computing (ASC) Program and the DOE Office of Science ASCR Program, and the Laboratory Directed Research and Development program (LDRD) at the Los Alamos National Laboratory.

REFERENCES

- [1] Bailey DS. Second-order monotonic advection in LASNEX. *Laser Program Annual Report '84*. UCRL-50021-84, Lawrence Livermore National Laboratory, 1984; 3-57-3-61.
- [2] Benson DJ. Computational methods in Lagrangian and Eulerian hydrocodes. *Computer Methods in Applied Mechanics and Engineering* 1992; **99**:235-394.
- [3] Pernice M, Walker HF. NITSOL: A Newton iterative solver for nonlinear systems. *SIAM Journal on Scientific Computing* 1998; **19**:302-318.

Barry Lee
**Improved Multiple-Coarsening/Semi-Coarsening Methods for S_n
Discretizations of the Boltzmann Equation**

Center for Applied Scientific Computing
L-561
Lawrence Livermore National Lab
Livermore
CA 94551
lee123@llnl.gov

In a recent series of articles, the author presented a multiple-coarsening multigrid method for solving S_n discretizations of the Boltzmann transport equation. This algorithm is applied to an integral equation for the scalar flux or moments. Although this algorithm is very efficient over parameter regimes that describe realistic neutron/photon transport applications, improved methods that can reduce the computational cost are presented in this talk. These improved methods are derived through a careful examination of the frequencies, particularly the near-nullspace, of the integral equation. In the earlier articles, the near-nullspace components were shown to be smooth in angle in the sense that the angular fluxes generated by these components are smooth in angle. In this talk, we present a spatial description of these near-nullspace components. Using the angular description of the earlier papers together with the spatial description reveals the intrinsic space-angle dependence of the integral equation's frequencies. This space-angle dependence then is used to determine the appropriate space-angle grids to represent and efficiently attenuate the near-nullspace error components on. It will be shown that these components can have multiple spatial scales. By using only the appropriate space-angle grids that can represent these spatial scales in the original multiple-coarsening algorithm, an improved algorithm is obtained. Moreover, particularly for anisotropic scattering, recognizing the strong angle dependence of the angular fluxes generated by the high frequencies of the integral equation, an improved multiple-coarsening scheme is derived. Restricting this scheme to the appropriate space-angle grids produces a very efficient method.

Si Liu
**Multilevel Domain Decomposition Algorithms for Inverse Elliptic
Problems**

2300 Arapahoe Ave
Apt 324
Boulder
CO 80302
sliu@colorado.edu
Xiao-Chuan Cai

Department of Computer Science, University of Colorado at Boulder

We investigate scalable multilevel domain decomposition algorithms for solving inverse elliptic problems formulated as optimization problems constrained by partial differential equations. To solve these optimization problems, we employ a fully coupled Lagrange-Newton-Krylov-Schwarz algorithm. One of the key steps in the algorithm is the Jacobian preconditioning, for which we study and compare four types of two-level domain decomposition methods. Our numerical results show that the algorithms work well for different types of observations in terms of the accuracy of the solution, and some of the algorithms scale better than the others when the number of processors is large. We also

study and report the sensitivity of the algorithms with respect to the jumps of the coefficients, the level of noise in the observed data, the size of the computational domain, the size of the mesh, and the number of processors.

Ira Livshits
**Algebraic multigrid solver for finding a full eigenbasis of the
two-dimensional Scghrodinger operator**

Department of Mathematical Sciences
Ball State University
Muncie
IN 47304
`ilivshits@bsu.edu`

A new multigrid solver for finding many eigen functions of the 2D Schrodinger operator is discussed. The solver employs multiscale eigenbasis whcih means an accurate representation of the finest-grid problem by an increasingly larger set of increasingly smaller eigen problems as the solver proceeds to coarser grids. Finally, on the coarsest grid many eigen problems, each of size $O(1)$, solved, collectively producing accurate approximations to all finest-grid eige pairs. The algorithm is built in an adaptive algebraic framework, using Brandt's least squares approach for prolongation operators, and Galerkin method for building coarse-grid operators.

Scott MacLachlan
**An angular multigrid method for modeling charged-particle transport in
Flatland**

Department of Mathematics
Tufts University
Bromfield-Pearson Hall
503 Boston Ave
Medford
MA 02155
`scott.maclachlan@tufts.edu`
Christoph Börgers
Department of Mathematics, Tufts University

Beams of microscopic particles penetrating scattering background matter play an important role in several applications. In this work, we consider parameter choices that are motivated by the problem of electron-beam cancer therapy planning. Mathematically, a steady particle beam penetrating matter, or a configuration of several such beams, is modeled by a boundary value problem for a Boltzmann equation. Grid-based discretization of such a problem leads to a system of algebraic equations, which is typically very large because of the large number of independent variables in the Boltzmann equation (six if no dimension-reducing assumptions other than time independence are made). If grid-based methods are to be practical for these problems, it is necessary to develop fast solvers for the discretized problems.

For beams of mono-energetic particles interacting with a passive background, but not with each other, in two space dimensions, an angular domain decomposition was proposed by Börgers in 1997. In this talk, we discuss an angular multigrid algorithm for the same model problem, based on a careful

choice of relaxation and coarse-grid correction processes. Our numerical experiments show rapid, grid-independent convergence for the forward-peaked scattering typical of electron beams. Unlike angular domain decomposition, the angular multigrid method works well even when the angular diffusion coefficient is fairly large.

Bram Metsch
Algebraic Multigrid (AMG) - Beyond the M-Matrix case

Universitaet Bonn
Institut fuer Numerische Simulation
Wegelerstrasse 6
53115 Bonn
GERMANY
`metsch@ins.uni-bonn.de`
Michael Griebel
Universitaet Bonn Institut fuer Numerische Simulation
Marc Alexander Schweitzer
Universitaet Bonn Institut fuer Numerische Simulation

Algebraic Multigrid methods (AMG) are known to be robust linear solvers for a wide class of linear systems arising from the discretization of partial differential equations. However, the construction as well as the convergence analysis of classical AMG methods is based on the assumption that the linear system matrix is a symmetric positive definite M -matrix.

In this talk we present an approach to extend the applicability of AMG to a wider class of problems. In particular, we focus on higher-order prolongation schemes to obtain a more accurate representation of the smooth error components.

William F Mitchell
The hp -Multigrid Method Applied to hp -Adaptive Finite Elements

100 Bureau Dr Stop 8910
NIST
Gaithersburg
MD 20899-8910
`william.mitchell@nist.gov`

Recently the hp version of the finite element method has received increasing attention. This is an adaptive finite element approach in which adaptivity occurs in both the size, h , of the elements (spacial or h adaptivity) and in the order, p , of the approximating piecewise polynomials (order or p adaptivity). The objective is to determine a distribution of h and p that minimizes the error using the least amount of work in some measure. The main attraction of hp adaptivity is that, in theory, the discretization error can decrease exponentially with the number of degrees of freedom, n , even in the presence of singularities, and this rate of convergence has been observed in practice.

It is desirable to combine this optimal order discretization method with an optimal order algebraic solution method, such as multigrid. Most published work in hp adaptive methods to date has focused on domain decomposition and preconditioned conjugate gradient type solvers or used direct

solvers. In this presentation we present the hp -multigrid method for high order finite elements and hp -adaptive grids.

An intriguing notion is to use the values of p as the levels of a multilevel method. This was first proposed in 1985 for the p -version of the finite element method by Craig and Zienkiewicz [1]. Several other researchers developed multilevel methods based on p in the later half of the 1980's, with Babuška et al. [2] coining the term multi- p (as opposed to multigrid) method. Theoretical p -independent (or nearly independent) convergence was established in the 1990's by Maday, Muñoz, Patera and Rønquist [3] and by Hu, Guo and Katz [4]. In the later 1990's and early 2000's the term p -multigrid was adopted, and it was mostly used in conjunction with the discontinuous Galerkin method for hyperbolic conservation laws.

The basic idea of the p -multigrid method is to use the p -hierarchical basis for the finite element space and perform a classic V-cycle in which the k^{th} level consists of all the bases up to degree k . The system of equations on level 1, i.e. the linear basis functions, is solved with a direct method. Nastase and Mavriplis [5] replaced the direct solver with a classic h -multigrid method, and coined the term hp -multigrid method.

Most of the results presented with p - and hp -multigrid methods has been with spectral element methods and discontinuous Galerkin methods. Although some of the papers mention hp -adaptivity, to the author's knowledge no papers have presented results using hp -adaptively refined meshes. In this talk we define the hp -multigrid method in the context of hp -adaptive finite elements for bisected triangular elements in 2D, and present numerical convergence results using elliptic problems with singular and nearly singular solutions. The results exhibit residual contraction factors that are independent of both h and p .

REFERENCES

- [1] Craig, A.W. and Zienkiewicz, O.C., A multigrid algorithm using a hierarchical finite element basis, in Multigrid Methods for Integral and Differential Equations, D.J. Paddon and Holstein (eds.), Clarendon Press, Oxford, 1985, pp. 310-312.
- [2] Babuška, I., Griebel, M. and Pitkaranta, J., The problem of selecting the shape functions for a p -type finite element, Internat. J. Numer. Methods Engng., 28, (1989), pp. 1891-1908.
- [3] Maday, Y., Muñoz, R., Patera, A.T. and Rønquist, E.M., Spectral element multigrid methods, in Proc. of the IMACS Int. Symposium on Iterative Methods in Linear Algebra, P. de Groen and R. Beauwens (eds.), Elsevier, Amsterdam, 1992, pp. 191-201.
- [4] Hu, N., Guo, X.-Z. and Katz, I., Multi- p preconditioners, SIAM J. Sci. Comput., 18, (1997), pp. 1676-1697.
- [5] Nastase, C.R. and Mavriplis, D.J., High-order discontinuous Galerkin methods using an hp -multigrid approach, J. Comput. Phys., 213, (2006), pp. 330-357.

David Moulton
Challenges for Multilevel Sampling in Statistical Inference

Applied Mathematics and Plasma Physics
MS B284

Los Alamos National Laboratory
Los Alamos
NM 87545

moulton@lanl.gov

Dave Higdon

Statistical Sciences, Los Alamos National Laboratory

Jasper Vrugt

Center for Nonlinear Studies, Los Alamos National Laboratory

Statistical inference through the Bayesian framework provides a powerful probabilistic approach to characterizing the solution of inverse problems. A canonical example is the recovery of the electrical conductivity of an object from measurements of potential and flux at the boundary, often called electrical impedance tomography (EIT). The mathematical model that maps the state (e.g., conductivity) to the measured data (e.g., potential and flux) is referred to as the forward model. In this statistical framework, solving the inverse problem corresponds to quantifying statistics about the posterior distribution of the state (e.g., conductivities) conditioned on measurements. Although this formulation is very flexible, the calculation of integrals with respect to the posterior distribution is typically done through Markov Chain Monte Carlo (MCMC) sampling, and hence it is a significant computational burden for complex multiscale forward models. To alleviate this burden applications often restrict their analysis to oversimplified forward models, in turn limiting the utility of their analysis.

To avoid these undesirable simplifications and additional uncertainty, considerable research has focused on schemes that improve the efficiency of the MCMC work horse, single-site Metropolis. In fact, a variety of approaches have been proposed to improve the efficiency of this rudimentary MCMC sampling scheme for high-dimensional posteriors. Multivariate updating approaches have been introduced that adjust multiple parameters simultaneously to construct a new proposal, while two-level delayed acceptance schemes rigorously allow the use of an approximate forward map (or approximate solve) in an initial screening step. In addition, high-level or templated-based priors (i.e., distributions of blobs or layers of conductivity) have been used to effectively reduce the dimensionality of the problem. Each of these techniques have demonstrated significant gains in efficiency for some problems, but achieving a significant improvement in scalability over a broad class of inverse problems seems to remain tied to reducing the dimensionality.

This desire to reduce the dimensionality creates an allure to leverage concepts from multilevel iterative solvers, which achieve their optimal scaling through the recursive use of coarser grids, with MCMC sampling. In this talk we highlight key issues in pursuing this connection, such as the interpretation of the hierarchy of coarse-scale models, as well as the scale-dependent interpretation of the model parameters and their distributions. Specifically, we explore the use of multilevel solvers within a delayed acceptance scheme, and discuss the potential speedup of this approach. Then to combine the strengths of this approach with concepts from high-level proposals we explore a framework that directly samples the discrete hierarchy of models.

Anna Naumovich
Defect correction with algebraic multigrid for the linearized steady Euler equations

German Aerospace Center
Institute of Aerodynamics and Flow Technology
Lilienthalplatz 7
D-38108 Braunschweig
Germany
anna.naumovich@dlr.de

Malte Foerster
Fraunhofer Institute for Algorithms and Scientific Computing, Germany
Richard Dwight

German Aerospace Center, Institute of Aerodynamics and Flow Technology, Germany

In the present work the linearized discrete compressible Euler equations are investigated. These equations are of importance for evaluation of solution derivatives, *a posteriori* error estimation, and frequency-domain methods for periodic flows. Moreover, their solution can be considered as a preliminary step towards developing algorithms for the non-linear equations. We focus on solution of 2nd-order accurate finite volume discretizations built on unstructured grids with either upwind or central convective fluxes. The goal is to reduce the CPU time required to obtain steady-state convergence for these given discretizations.

Nowadays, such problems are typically solved with geometric multigrid methods, employing either explicit Runge-Kutta or an approximate implicit scheme as a smoother. Despite the investigation of many variations on this theme over the years there has been a continued lack of success in obtaining satisfactory convergence for all but the simplest geometries. The situation is even worse for unstructured discretizations, where the additional complexity of definition of a proper coarse grid hierarchy arises.

In this work defect correction [1, 2] is applied to solve the linear problem. Defect correction allows by-passing direct solution of the 2nd-order accurate discretization, instead, a sequence of problems with a 1st-order discretization on the left-hand side is solved. On this better conditioned problems we apply AMG.

Since algebraic multigrid [4, 5, 6] does not rely on the geometry of the grid, it is an attractive alternative to geometric multigrid for problems discretized on unstructured grids. However, initially developed for scalar elliptic PDEs, algebraic multigrid needs special extensions to be applicable to systems of PDEs, one of those extensions being a point-based approach. The point-based AMG solver, implemented in the SAMG package [3], was investigated in this work. The solver, applied together with accelerator BiCGStab, proved to be very efficient for solution of the 1st-order accurate discretizations considered here in sub-, tran- and super-sonic flow regimes. One should mention that our attempts to apply AMG directly to 2nd-order discretizations were not much successful so far.

The defect correction approach combined with the AMG solver was applied to solve 2nd-order accurate discretizations and was seen to significantly out-perform the geometric multigrid solver [7, 8] in terms of CPU time for all flow regimes. This is mainly by virtue of the efficiency of AMG, which requires very few iterations to give sufficient accuracy of the inner problem at each defect correction step. Furthermore, as the 1st-order matrix is constant, only one AMG setup phase must be performed for the entire calculation.

REFERENCES

- [1] K. Boehmer, P. Hemker, B. Van Leer, The defect correction approach, *Computing Suppl.*, 5, 1984.
- [2] B. Koren, Multigrid and defect correction for the steady Navier-Stokes equations, *J. Comput. Phys.* (ISSN 0021-9991), vol. 87, 1990, pp. 25-46.
- [3] K. Stueben, T. Clees, SAMG Users Manual, Fraunhofer Institute SCAI, <http://www.scai.fraunhofer.de/samg>.
- [4] A. Brandt, S.F. McCormick, J. Ruge, Algebraic multigrid (AMG) for sparse matrix equations, *Sparsity and its Applications*, D.J. Evans (ed.), Cambridge University Press, Cambridge, 1984, pp. 257-284.
- [5] J.W. Ruge, K. Stueben, Algebraic Multigrid (AMG) , in *Multigrid Methods* (McCormick, S.F., ed.), SIAM, *Frontiers in Appl. Math.*, Vol. 5, Philadelphia, 1986.
- [6] K. Stueben, Algebraic Multigrid (AMG): An Introduction with Applications, GMD Report 70, Nov. 1999.
- [7] D. Schwamborn, T. Gerhold, R. Heinrich, The DLR TAU-Code: Recent Applications in Research and Industry, in P. Wesseling, E. Onate and J. Periaux (eds) *European Conference on Computational Fluid Dynamics. ECCOMAS CFD 2006*.
- [8] R. Dwight, J. Brezillon Efficient and Robust Algorithms for Solution of the Adjoint Compressible Navier-Stokes Equations with Applications, *Int. J. Num. Meth. Fluids*, DOI: 10.1002/fld.1894, 2008.

Pablo Navarrete
**Signal Processing Approach to avoid Smoothing Iterations in Multi-grid
Methods**

Department of Electrical Engineering
Universidad de Chile
Av Tupper 2007
Santiago
RM 8370451
Chile
pnavarrete@ing.uchile.cl

Modifications of the conventional multi-grid algorithm are studied to avoid the use of smoothing iterations. In the full multigrid algorithm, classical smoothing iterations (e.g. Gauss-Seidel) reduce high-frequency components of the error and a coarse-grid approach reduces the low-frequency components of the error. The problem here is that two methods with different structures are being combined, which introduces additional complexity in the convergence analysis of multi-grid methods. Then, the idea is to avoid the use of smoothing iterations by using different inter-grid configurations and the concept of quadrature mirror filters, which are well known in the area of signal processing and particularly in wavelet analysis. This framework can be introduced by using the structure of the extended convergence analysis introduced in [1] from which the classical Local Fourier Analysis (LFA) is a particular case. This can provide an integrated configuration to efficiently reduce low- and high- frequency components of the error as well as aliasing effects between the low- and high-frequency components of the error. The possibility of a direct solver is studied and the conditions under which it can be implemented.

REFERENCES

- [1] P. Navarrete, and E.J. Coyle, “A semi-algebraic approach that enables the design of inter-grid operators to optimize multigrid convergence,” *Numerical Linear Algebra with Applications*, Vol. 15, No. 2-3, pp. 219-247, March 2008.

Yvan Notay
Algebraic analysis of MG methods: the nonsymmetric case

Universite Libre de Bruxelles
Service de Metrologie Nucleaire (CP 165/84)
50 av F D Roosevelt
1050 Bruxelles
BELGIUM
ynotay@ulb.ac.be

We develop an algebraic analysis of two-grid methods that does not require any symmetry property. Several equivalent expressions are provided that characterize *all* eigenvalues of the iteration matrix. In the symmetric positive definite case, these expressions reproduce the sharp two-grid convergence estimate obtained by Falgout, Vassilevski and Zikatanov [*Numer. Lin. Alg. Appl.*, 12 (2005), pp. 471–494], and also previous algebraic bounds which can be seen as corollaries of this estimate.

These results allow to measure the convergence by checking “approximation properties”. In this talk, proper extensions of the latter to the nonsymmetric case are presented. Sometimes approximation

properties for the symmetric positive definite case are summarized in loose terms; e.g.: *Interpolation must be able to approximate an eigenvector with error bound proportional to the size of the eigenvalue* [SIAM J. Sci. Comput., 22 (2000), pp. 1570–1592]. It is shown that this can be applied to nonsymmetric problems too, understanding “size” as “modulus”.

Eventually, an analysis is developed, for the nonsymmetric case, of the theoretical foundations of “compatible relaxation”, according to which a Fine/Coarse partitioning may be checked and possibly improved.

Kees Oosterlee
Numerical Mathematics Aspects in Computational Finance

Delft University of Technology
Delft Institute of Applied Mathematics (DIAM)
Mekelweg 4
2628 CD Delft
The Netherlands
C.W.Oosterlee@cwi.nl

In this presentation we will discuss some topics in Finance that require Mathematics, and, in particular, efficient numerical techniques. We will discuss option pricing, for example, for option contracts based on more than one underlying stock, and contracts with advanced stochastic models for the underlying stock price dynamics. We will focus on a mathematical framework in which we perform this research. One of the aims is to price the options as fast as possible.

Minho Park
A NEW LEAST SQUARES BASED AMG

1300 30TH ST APT E1-14
Boulder
CO 80303
parkmh@colorado.edu
Steve McCormick
University of Colorado
Tom Manteuffel
University of Colorado
John Ruge
University of Colorado
Marian Brezina
University of Colorado

In this talk we discuss the performance of new least squares based algebraic multigrid method which is modification of Brandt’s Bootstrap AMG(BAMG). Solving the linear system using least squares based AMG requires set of vectors that are results of several fine level relaxation sweeps on homogeneous equation $Ax = 0$. Unlike Brandt’s BAMG, new method approximate all F-F connections in least squares sense. The presented numerical experiments demonstrate that the method can achieve good convergence with less vectors and relaxation sweeps.

Michael Pernice
Multilevel solution of diffusion operators with cut cell discretizations

Center for Advanced Modeling and Simulation
Idaho National Laboratory
P O Box 1625
Idaho Falls
ID 83415-3550
michael.pernice@inl.gov

Diffusion operators $\mathcal{L} = \nabla \cdot \beta \nabla$ with discontinuous coefficients $\beta : \mathbb{R}^d \rightarrow \mathbb{R}$ are encountered in numerous applications, such as heat conduction, neutron transport, and subsurface flow. These discontinuities represent changes in material properties (thermal conductivity, interaction cross sections, and permeability). Accurate discretization of these operators is challenging when the discontinuities in β are not aligned with the computational mesh. Mesh alignment can be achieved by constraining mesh generation to conform to the geometry of the different materials when using unstructured or boundary-fitted Cartesian grids. However for complicated geometries this may be intractable. An alternative approach is to use a background Cartesian grid and to represent the geometry by intersecting it against the grid. This gives rise to a set of irregular *cut cells* where Cartesian grid cells intersect the geometry.

While originally developed to handle irregular physical boundaries, cut cells can also be used to resolve material interfaces that are interior to the domain. This approach replaces the difficult global grid generation problem with more tractable local intersection problems. When combined with local mesh refinement, high resolution representation of complex geometry can readily be obtained. While this essentially solves the geometry problem, satisfactory discretization of diffusion operators has remained elusive. Second-order finite difference schemes that account for the location and orientation of the interface in a cut cell have been developed, but these approaches are not conservative. Finite volume schemes, such as the ghost fluid method developed in the context of interface tracking using level sets, have also been developed. However analysis of the convergence of these methods is complicated by the lack of smoothness of the solution in the vicinity of the jumps in β .

Because local mesh refinement is necessary for accurate and efficient representation of complex geometry with cut cells, multilevel solvers are essential. We discuss issues related to discretization, grid convergence, and solver efficiency for diffusion operators on cut cell grids.

Xavier PINEL
**Theoretical and experimental analysis of a perturbed two-grid
preconditioner for indefinite Helmholtz problems**

CERFACS
42 Avenue Coriolis
31057 TOULOUSE cedex 1
FRANCE
`xavier.pinel@cerfacs.fr`
Henri CALANDRA
*TOTAL, Centre Scientifique et Technique Jean Feger, Avenue Larribau, F-64018 Pau cedex,
France, Henri.CALANDRA@total.com*
Serge GRATTON
*CNES, avenue Edouard Belin, F-31401 Toulouse cedex 4, France and CERFACS, 42 avenue
Gaspard Coriolis, F-31057 Toulouse cedex 1, France, serge.gratton@cnes.fr and
Serge.Gratton@cerfacs.fr*
Xavier VASSEUR
*CERFACS, 42 avenue Gaspard Coriolis, F-31057 Toulouse cedex 1, France,
Xavier.Vasseur@cerfacs.fr*

We study the three-dimensional Helmholtz equation written in the frequency domain modeled by the following partial differential equation:

$$-\Delta u - k^2 u = g$$

with some absorbing boundary conditions, where u is the pressure of the wave, k its wavenumber and g is a Dirac function that represents the wave source. This problem is discretized using second-order accurate finite difference techniques leading to huge linear systems for large wavenumbers.

Following Elman [1], we use a geometric two-grid preconditioner for a Krylov subspace method (namely flexible GMRES [2]) as a solution method. Due to the large dimension of the coarse grid problem we focus on the behavior of a two grid algorithm where the coarse problem is not solved exactly. We use a Krylov subspace method to solve this problem only approximately. This leads us to analyze two questions:

- Which stopping criterion and convergence threshold should be used for the coarse grid solver ?
- What is the impact of an approximate coarse grid solution on the preconditioning properties ?

In this talk, we will bring some elements to answer both questions. First a local Fourier analysis - assuming Dirichlet boundary conditions and small wavenumbers - will provide convergence rate estimates for this perturbed two-grid algorithm used as a solver. Secondly we will investigate the numerical behaviour of the algorithm when this two-grid method is used as a preconditioner. For that purpose we study the spectrum of an equivalent matrix constructed inside the flexible variant of GMRES. This will help us understanding the effects of the approximate coarse grid solution on the preconditioner for both absorbing boundary conditions and large wavenumbers. Parallel numerical experiments will conclude this talk showing the efficiency of this perturbed preconditioner even for large wavenumbers.

REFERENCES

- [1] Elman, H. R., Ernst, O. G. and O’Leary, D. P., A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations. *SIAM J. Sci. Comput.*, **23**, 1291–1315, 2001.
- [2] Saad, Y., A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Sci. Comput.*, **14**, 461–469, 1993.

Ulrich Ruede
How fast are multigrid methods?

Universitaet Erlangen
Lehrstuhl fuer Simulation
Cauerstrasse 6
91058 Erlangen
Germany
ruede@cs.fau.de

The most efficient multigrid algorithms known require less than 30 operations per unknown to solve the 5-point stencil discretization of the Poisson equation in 2D to the level of truncation error. This uses full multigrid and a clever implementation of an FAS-like scheme using hierarchical basis transformations. We will briefly review the features of this algorithm will also discuss its generalizations for other problems. Unfortunately, the real performance of multigrid algorithms on current computer systems is much lower than predicted by such work estimates. We will discuss why this is the case, and what can be done about it.

Martin Rupp
Eigenmodes of a Guitar Top Plate – Application for Filtering Algebraic Multigrid and Preconditioned Inverse Iteration

Simulation and Modelling
Goethe-Center for Scientific Computing
Goethe-University
Kettenhofweg 139
60325 Frankfurt am Main
Germany
martin.rupp@gcsc.uni-frankfurt.de
Arne Nägel
Goethe-University, Frankfurt/Main, Germany
Gabriel Wittum
Goethe-University, Frankfurt/Main, Germany

The top plate of a guitar is crucial for sound amplification. The frequencies in the spectrum of the vibrating string which are close to resonance frequencies of the top plate are amplified the most, and determine (among other factors) the sound of the guitar. We employ a finite element discretisation of the Lamé equation for the numerical simulation of this problem.

To compute the resonance frequencies, we use a projection method for generalized symmetric eigenvalue problems $Au = \lambda Mu$. The key step is a minimization of the Rayleigh-Quotient $\lambda(u) := \frac{\langle Au, u \rangle}{\langle Mu, u \rangle}$

over a suitable constructed subspace \mathcal{L} via the Rayleigh-Ritz-method. To construct \mathcal{L} , we use the *Preconditioned Inverse Iteration* (PINVIT): The approximate solution of the Inverse Iteration equation leads to the correction

$$v^{(k+1)} = v^{(k)} - c^{(k)} = v^{(k)} - B(A - \lambda(v^{(k)})M)v^{(k)} \quad (8)$$

of the eigenvalue approximation $v^{(k)}$, which motivates the use of

$$\mathcal{L}^{(k+1)} = \text{span} \langle c^{(k)}, v^{(k)} \rangle .$$

Additional vectors in \mathcal{L} may increase convergence like in *LOPCG* [1] or, more general, in PINVIT(s) [2].

The preconditioner B used in (8) should approximate the inverse of the stiffness matrix A , and must thus be robust with respect to anisotropies in the geometry and the material properties. We seek to treat the arising linear system using algebraic multigrid. In particular, we focus on the *Filtering Algebraic Multigrid* approach [3,4]. The key idea of this AMG variant is to construct the interpolation operator P , such that the norm of the two-grid operator is minimized in a certain sense. At the same time, constraints are imposed to guarantee filter conditions for certain test vectors t :

$$\min_P \|(I - PR^{(inj)})S\|, \quad \text{s.t. } (I - PR^{(inj)})St = 0$$

We comment on the theory and outline a pointwise version of the method, which is suitable to treat systems of equations. For linear elasticity problems, local representations of the rigid body modes are used to obtain robustness on the Neumann boundaries. Numerical results and illustrative examples are provided and conclude the talk.

REFERENCES

- [1] Knyazev, A.V. : *Toward the Optimal Preconditioned Eigensolver: Locally Optimal Block Preconditioned Conjugate Gradient Method*. SIAM J. Sci. Comput., Vol. 23, No. 2, pp. 517–541, 2001.
- [2] K. Neymeyr: *On preconditioned eigensolvers and Invert-Lanczos processes*. Preprint 2007.
- [3] Wagner C.: *On the algebraic construction of multilevel transfer operators*. Computing, Vol. 65, No. 1, pp. 73–95, 2000.
- [4] Nägel A., Falgout R.D., Wittum G.: *Filtering algebraic multigrid and adaptive strategies*. Computing and Visualization in Science, Vol. 11, No. 3, pp. 159–167, 2008.

Geoffrey D. Sanders
Acceleration of Algebraic Multigrid Methods for Markov Chains

Department of Applied Mathematics
526 UCB
University of Colorado
Boulder
CO 80309-0526
sandersg@colorado.edu
Hans De Sterck
University of Waterloo
Tom Manteuffel
University of Colorado, Boulder
Steve McCormick
University of Colorado, Boulder
Killian Miller
University of Waterloo
John Ruge
University of Colorado, Boulder

In many application areas, including web page ranking and networking systems, finding the steady-state distribution vector of a Markov system is of interest, and often difficult to compute efficiently. The steady-state vector is the solution to a nonsymmetric eigenproblem $B\mathbf{x} = \mathbf{x}$, subject to probability constraint $\|\mathbf{x}\|_1 = 1$, where B is column stochastic, $\mathbf{1}^t B = \mathbf{1}^t$. Recently, scalable methods involving Smoothed Aggregation (SA) and Algebraic Multigrid (AMG) were proposed to solve such eigenvalue problems. These methods use multiplicative iterate updates versus additive error corrections that are typically used in linear solvers. This work discusses the implementation of an outer iteration to the existing methods that accelerates convergence of multiplicative update methods, similar in principle to a preconditioned flexible Krylov wrapper applied to a linear problem. The acceleration is performed by selecting a linear combination of old iterates to minimize a functional that has the steady-state solution direction as a unique minimizing zero. The quality of the acceleration is demonstrated with a few simple examples.

Jacob Schroder
Smoothed Aggregation Multigrid for One-Dimensional Helmholtz Problems

Siebel Center for Computer Science
4328 Siebel Center
201 N Goodwin Ave
Urbana
IL 61801
USA
jacob.bb.schroder@gmail.com
Luke Olson

Assistant Professor, Computer Science Department, University of Illinois at Urbana-Champaign

We present a purely algebraic smoothed aggregation (SA) multigrid method for one-dimensional scalar Helmholtz problems with exterior radiation boundary conditions. The scalar Helmholtz problem is particularly difficult for algebraic multigrid (AMG) solvers. Not only can the discrete operator

be indefinite and non-self-adjoint, but it also allows for oscillatory error components that yield relatively small residuals. These oscillatory error components are not effectively handled by either standard relaxation or standard coarsening procedures. We address these difficulties through modifications of the SA method and by providing the setup phase with appropriate wave-like near-nullspace candidates. Our results for GMRES preconditioned with the proposed SA method exhibit near grid-independent performance, along with consistent performance across a wide range of frequencies.

Benjamin Seibold
AMG for Meshfree Finite Difference Methods

Department of Mathematics
2-346
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge MA 02139
USA
`seibold@math.mit.edu`

Lagrangian particle methods, such as Smoothed Particle Hydrodynamics, approximate the equations of fluid flow on a cloud of points that move with the flow. Thus, convective terms are treated exactly, which is a great advantage compared to Eulerian approaches. The fundamental challenge in particle methods are unstructured node distributions. A common strategy is to approximate differential operators by meshfree finite difference stencils. These can be derived for instance by the moving least squares method.

Incompressible flows can be simulated by a pressure correction in each time step. The arising Poisson equations can be approximated on the cloud of particles by meshfree finite differences. The generation and solution of the arising linear systems is often times very costly, thus multigrid methods are of great interest. A natural approach is Algebraic Multigrid, since geometric strategies to refine a fully unstructured point cloud are difficult to formulate.

The finite difference matrices arising from classical meshfree least squares approaches have three properties that are challenging with respect to AMG:

- Finite difference methods on unstructured data sets approximate the Laplace operator by non-symmetric matrices.
- The arising matrices are about six times as dense as finite difference matrices on rectangular grids.
- An M-matrix structure is not guaranteed, and is—unless the point geometry is sufficiently nice—in general violated.

While here the non-symmetry is in some sense weaker than coming from the approximation of a convection operator, it can not be overcome. Better sparsity and the M-matrix structure, however, *can* be achieved, by a new finite difference approach, that is based on linear sign-constrained optimization.

In this talk, I will outline meshfree finite difference methods, present the new finite difference approach, show results when applied in the context of AMG methods, and pose open questions. Numerical tests were done using SAMG (Fraunhofer SCAI, K. Stüben et al.), as well as AMLI (C. Mense and R. Nabben).

Chris Siefert
**An Algebraic Multigrid Method for Least-Squares Finite Element
Systems**

Sandia National Laboratories
P O Box 5800
MS 0378
Albuquerque
NM 87185-0378
csiefer@sandia.gov
Pavel Bochev
Sandia National Laboratories
Kara Peterson
Sandia National Laboratories

Least-squares finite element methods are an attractive class of methods for the numerical solution of certain partial differential equations, specifically div-curl systems. They are motivated by the desire to recover, in more general settings, the advantageous features of Rayleigh-Ritz methods such as the avoidance of discrete compatibility conditions and the production of symmetric and positive definite discrete systems. The methods are based on the minimization of convex functionals that are constructed from equation residuals.

This work deals with formulation of a class of compatible least-squares principles for a model div-curl equation and the efficient multigrid solution of the resulting algebraic systems. Specifically, we consider a discrete least-squares principle

$$\min_{u^h \in C_0^h} \|\nabla \times u^h - g\|_0^2 + \|\nabla_h^* \cdot u^h - f\|_0^2 \quad (9)$$

where C_0^h is a curl-conforming space constrained by the homogeneous boundary condition $u \times \mathbf{n} = 0$ on $\partial\Omega$, and ∇_h^* is a discrete divergence operator acting on that space.

The use of discrete divergence in this formulation is compelled by the well-known problems in the finite element approximation of the Hilbert space $H(\text{curl}, \Omega) \cap H(\text{div}, \Omega)$. In particular, in (9) we approximate that space conformally with respect to the curl operator. This yields a linear system in the form of a Hodge Laplacian discretized using edge (curl-conforming) elements. We demonstrate that this system can be solved using a multigrid method derived from the previous work of Bochev, Hu, Siefert and Tuminaro [1]. A complementary least-squares formulation wherein $H(\text{curl}, \Omega) \cap H(\text{div}, \Omega)$ is approximated by div-conforming elements yields a linear system corresponding to a Hodge Laplacian discretized by face (div-conforming) elements. To solve this system we use an alternative formulation of the multigrid method derived in [2].

REFERENCES

- [1] P. Bochev, J. Hu, C. Siefert, and R. Tuminaro. An algebraic multigrid approach based on a compatible gauge reformulation of Maxwell's equations. *SIAM J. Sci. Comp.* Volume 31, Issue 1, pp. 557-0583, 2008.
- [2] P. Bochev, C. Siefert, R. Tuminaro, J. Xu, and Y. Zhu, Compatible Gauge Approaches for $H(\text{div}, \Omega)$ Equations. CSRI Summer proceedings, Sandia National Laboratories Report SAND-7977, pp. 3-14, 2007.

Georg Stadler
Experience with parallel AMG for large-scale adaptive mantle convection

Institute for Computational Engineering and Sciences (ICES)
The University of Texas at Austin
1 University Station
C0200
Austin
TX 78712
`georgst@ices.utexas.edu`
Carsten Burstedde

*Institute for Computational Engineering and Sciences (ICES) The University of Texas at Austin,
Austin, TX*

Omar Ghattas

*Institute for Computational Engineering and Sciences (ICES) and Depts. of Geological Sciences
and Mechanical Engineering The University of Texas at Austin, Austin, TX*

Michael Gurnis

Seismological Laboratory, California Institute of Technology, Pasadena, CA

Eh Tan

Computational Infrastructure for Geodynamics, California Institute of Technology, Pasadena, CA

Tiankai Tu

*Institute for Computational Engineering and Sciences (ICES) The University of Texas at Austin,
Austin, TX*

Lucas W. Wilcox

*Institute for Computational Engineering and Sciences (ICES) The University of Texas at Austin,
Austin, TX*

Zhong, Shijie

Department of Physics, University of Colorado, Boulder, CO

Mantle convection is the principal control on the thermal and geological evolution of the Earth. Its modeling involves solution of the mass, momentum, and energy equations for a viscous, creeping, incompressible non-Newtonian fluid at high Rayleigh and Peclet numbers. We are developing the parallel adaptive mantle convection code **Rhea**, which builds on our **Adaptive Large-scale Parallel Simulations** (ALPS) framework.

The numerical simulation of mantle convection requires the solution of a stationary, highly variable-viscosity Stokes system at each time step. The Stokes equation is discretized using trilinear finite elements for both velocity and pressure, and polynomial pressure projection is used to stabilize this equal-order approximation. For the solution of the discrete Stokes saddle point system we use the preconditioned minimal residual (MINRES) method. The preconditioner is based on an approximate Schur complement for the pressure component and involves a solve with a positive definite operator for the velocity field. This block solve is approximated with one V-cycle of algebraic multigrid. Results obtained with the parallel AMG libraries *BoomerAMG* from *hypre* and with *ML* from *Trilinos* will be shown. We discuss the parallel scalability of the solver on up to 16K cores, as well as its dependence on viscosity variations and boundary conditions.

Lei Tang
**Parallelization of efficiency-based adaptive local refinement for
FOSLS-AMG**

Department of Applied Mathematics
526 UCB

University of Colorado
Boulder
CO 80309-0526

`lei.tang@colorado.edu`

Marian Brezina

Department of Applied Mathematics, University of Colorado at Boulder
Tom Manteuffel

Department of Applied Mathematics, University of Colorado at Boulder
Steve McCormick

Department of Applied Mathematics, University of Colorado at Boulder
Josh Nolting

GeoEye Inc., Thornton, Colorado

John Ruge

Department of Applied Mathematics, University of Colorado at Boulder

A new adaptive local refinement strategy (ACE), in which the refinement decisions are based on optimizing the computational efficiency, has recently been developed. The first-order system least-squares (FOSLS) finite element methods, together with algebraic multigrid methods (AMG), complement this idea, resulting in a powerful tool for numerical solutions of many PDEs. However, in parallel environment, the global sorting of element local error indicators used by the original approach results in undesirable communication between processors. A strategy circumventing the global sorting is proposed. The test results show that the strategy has the potential of greatly reducing communication, while producing a sequence of refined grids with essentially the same quality as before.

Jari Toivanen
**An AMG preconditioner based on damped operators for time-harmonic
wave equations**

Institute for Computational and Mathematical Engineering

Durand Building

Room 023A

Stanford University

Stanford

CA 94305

`toivanen@stanford.edu`

Tuomas Airaksinen

University of Jyväskylä, Finland

Anssi Pennanen

University of Jyväskylä, Finland

An algebraic multigrid approximation of the inverse of the physically damped operators is used as a preconditioner for time-harmonic scattering problems in fluids and solids. The AMG uses graph based coarsening together with underrelaxed Jacobi smoother. Several numerical experiments demonstrate

the behavior of the method in complicated two-dimensional and three-dimensional domains. The number of iterations behaves roughly linearly with respect to the frequency. This approach leads to efficient solution procedure for low and medium frequency problems.

Panayot S. Vassilevski
Coarse Spaces by Constrained Energy Minimization

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
P.O. Box 808, L-560
Livermore, CA 94551, U.S.A.
panayot@llnl.gov

We consider an unified approach of constructing operator-dependent discretization spaces on relatively coarse computationally feasible meshes. The approach utilizes natural energy functionals associated with the PDEs of interest. We construct local basis functions by minimizing the underlined functional subject to a set of constraints. The constraints are chosen so that the resulting spaces possess increasingly high order of approximation. We investigate the proposed approach from an upscaling discretization point of view which we illustrate with some preliminary numerical examples.

Sergii Veremieiev
Multigrid strategies for the efficient and accurate solutions of free-surface film flow on man-made and naturally occurring functional substrates

School of Mechanical Engineering
University of Leeds
Leeds
West Yorkshire
United Kingdom
LS2 9JT
s.veremieiev@leeds.ac.uk
Phillip H. Gaskell

School of Mechanical Engineering, University of Leeds, Leeds, West Yorkshire, United Kingdom. LS2 9JT

Yeaw Chu Lee

School of Mechanical Engineering, University of Leeds, Leeds, West Yorkshire, United Kingdom. LS2 9JT

Harvey M. Thompson

School of Mechanical Engineering, University of Leeds, Leeds, West Yorkshire, United Kingdom. LS2 9JT

The deposition and flow of continuous thin liquid films over man-made or naturally occurring functional substrates containing regions of micro-scale topography (which may be fully submerged or extend through the free-surface of the film itself) plays an important role in numerous engineering and biologically related fields. For example, in the context of engineering processes thin film flows form a key role in photo-lithography [1] and precision coating processes [2], while in biological systems they occur in areas as diverse as tissue engineering [3] and plant disease control [4].

Current analytic and experimental methods are incapable, and likely to remain so for the foreseeable future, of either meeting the considerable challenges posed in unravelling the underlying physics

of so wide a range of real problems or of providing the necessary insight for improving man-made functional substrates and the design and creation of novel ones. Equally, it would be naive to pretend that there currently exists an off-the-shelf computational answer involving a combined multi-scale modelling approach comprised of a strategic mix of molecular dynamics, meso-scale and continuum methods. The reality is that at present, the modelling of three-dimensional free-surface film flows over substrates containing complex topography is still at an early stage of development, for which exploitation of the long-wave approximation is the focus. The key simplifying feature of the latter is that it enables the reduction of the governing time-dependent Navier-Stokes equations to a more tractable coupled system of partial differential equations; nevertheless the resulting reduced equation set is still required to be solved both efficiently and accurately.

In the context of the above, two such formulations are explored and solved for: (i) a simple lubrication (LUB) approach [5], where the dependent variables are the film height and pressure; (ii) a novel depth averaged form (DAF), which involves the determination of three variables - two velocity components and the film height. Approach (ii), unlike (i), enables thin film flows with inertia to be predicted and its effect quantified. Solution of the nonlinear problems of interest is extremely challenging since large computational domains and fine mesh structures are required: (a) to ensure grid independent solutions; (b) to capture persistent free surface disturbances caused by both localised (single) and distributed (multiply-connected) topography; (c) as the length-scale of the topographical features involved become smaller and the resolution required to capture the resultant flow accurately becomes increasingly important.

The research reported describes the development and application of an efficient, multigrid algorithm to solve the two different resulting equation sets, which implements the full approximation storage (FAS) and full multigrid (FMG) schemes together with a standard fixed number of pre- and post-smoothing V-cycles and appropriate inter-grid transfer and smoothing operators paired with Newton-Raphson iteration; appropriate account is also taken of the co-located (LUB) or staggered (DAF) meshing strategies employed. The associated time-descretisation includes the use of an explicit predictor [6] and a semi-implicit β -method [7] solution stages where automatic adaptive time-stepping is controlled in terms of the local truncation error on the finest grid level.

The utility of the multigrid methodology is explored in two different ways:

1. As a scalable parallel, portable object-oriented algorithm implemented on the following HPC architectures: HECToR, HPCx and BlueGene/P.
2. As a serial algorithm but with the additional feature of error controlled automatic mesh adaption [8].

In terms of achieving efficient and accurate solutions, the former is shown to deliver the expected benefits that parallelisation and the use of multiple processor computing platforms bring and to be particularly well suited to predicting thin film flow on surfaces containing densely packed and complex topographical features; the latter is found to be preferentially suited to the case of film flow over localised or sparsely distributed topographical features where mesh adaption leads to considerable savings in CPU times without loss of accuracy.

The examples chosen to illustrate the range of applicability of the above solution strategies are centred on the gravity-driven flow of thin films over substrates inclined at an angle θ to the horizontal and include: (i) comparison of the accuracy of predicted thin film profiles against experimental and complementary finite element solutions of the full Navier-Stokes problem, for varying flow parameters and both LUB and DAF approaches, in the case of two-dimensional span-wise topographies such as a step-up, a step-down, trenches and peaks; (ii) flow past three-dimensional localised surface patterns illustrating the influence of inertia and its effect with varying topography shape and size; (iii) the

analysis of flow over an engineered man-made multiple-connected surface pattern where the practical goal is that of minimising free-surface disturbances caused by the topography; (iv) thin film flow over leaf sections exhibiting heterogeneous surface patterning.

REFERENCES

- [1] W.K. Ho, A. Tay, L.L. Lee, C.D. Schaper, “On control of resist film uniformity in the microlithography process”, *Control. Eng. Prac.*, 12(7), 881-892, 2004.
- [2] A. Clarke, “Coating on a rough surface”, *A.I.Ch.E. Journal*, 48, 2149-2156, 2002.
- [3] D.P. Gaver, J.B. Grotberg, “The Dynamics of a Localized Surfactant on a Thin-Film”, *J. Fluid Mech.*, 213, 127-148, 1990.
- [4] D. R. Walters, ‘Disguising the leaf surface: the use of leaf coatings for plant disease control’, *Euro. J. Plant Pathology*, 114, 255-260, 2006.
- [5] P.H. Gaskell, P.K. Jimack, M. Sellier, H.M. Thompson, M.C.T. Wilson, “Gravity-driven flow of continuous thin liquid films on non-porous substrates with topography”, *J. Fluid Mech.*, 509, 253-280, 2004.
- [6] P.H. Gaskell, P.K. Jimack, M. Sellier, H.M. Thompson, “Efficient and accurate time adaptive multigrid simulations of droplet spreading”, *Int. J. Num. Meth. Fluids*, 45, 1161-1186, 2004.
- [7] T.J. Chung, *Computational Fluid Dynamics*. Cambridge: Cambridge University Press, 2002.
- [8] Y.C. Lee, H.M. Thompson, P.H. Gaskell, “An efficient adaptive multigrid algorithm for predicting thin film flow on surfaces containing localized topographic features”, *Computers & Fluids*, 38, 838-855, 2007.

Justin Wan
**A Multigrid Method for Solving Partial Integro-Differential Equations
 Arising from Option Pricing Under the Levy Process**

School of Computer Science
 University of Waterloo
 200 University Ave West
 Waterloo
 ON N2L 3G1
 Canada
 jwlwan@uwaterloo.ca
 Iris R. Wang
University of Waterloo

Recently, Levy process models for pricing options have become popular in the financial literature. Based on the Black-Scholes model, the option value of a contingent claim, $V(S, \tau)$, under the Levy process satisfies the following partial integro-differential equation (PIDE):

$$V_\tau = \frac{\sigma^2}{2} S^2 V_{SS} + (r - q) S V_S - rV + \int_{-\infty}^{\infty} \nu(y) [V(S e^y, \tau) - V(S, \tau) - S(e^y - 1) V_S] dy,$$

where S is the underlying stock price, τ the time to expiry, r the risk-free interest rate, q the dividend yield, and σ the volatility associated with the continuous component of Levy process. The Levy measure $\nu(y)$ is typically singular at $y = 0$.

Due to the singularity of $\nu(y)$, standard discretization methods may not achieve the usual accuracy. Moreover, the computation of the integral term can be very expensive. Recently, an implicit discretization method has been proposed that can obtain second order convergence for finite variation case and better than first order accuracy for infinite variation processes. In each time step, a linear system needs to be solved, which forms the bottle neck of the entire computation. A fixed point and a preconditioned BiCGSTAB iterations have been proposed. While they converge quite rapidly for mildly singular cases, the number of iterations grows significantly in the case of infinite variation when the mesh size decreases.

In this talk, we present a multigrid method for solving the PIDE. The multigrid method uses a fixed point iteration for smoothing. We prove by Fourier analysis that the smoother damps away the high frequencies. Linear interpolation and full weighting are used for intergrid transfer. Since it is too expensive to form the integral term explicitly, direct discretization is used for constructing the coarse grid matrices. However, it turns out that it is not as effective as the Galerkin course grid operators. We discuss how a combination of the two is used to form the coarse grid matrices. We demonstrate numerically the effectiveness of the multigrid method and show that the convergence rate is independent of the mesh size.

Roman Wienands
On the Construction of Prolongation Operators for Multigrid

Mathematical Institute
 University of Cologne
 Weyertal 86-90
 50931 Cologne
 Germany

`wienands@math.uni-koeln.de`

Harald Koestler

Department of Computer Science 10, University of Erlangen-Nuremberg, Germany

Irad Yavneh

Department of Computer Science, Technion-Israel Institute of Technology

The quality of multigrid methods crucially depends on an efficient interplay between the smoothing operator \mathbf{S} and prolongation \mathbf{P} . More precisely, the range of \mathbf{S} should be sufficiently well represented in the range of \mathbf{P} . In the limit case of an optimal (but usually impractical) choice of \mathbf{P} and \mathbf{S} this means that

$$\mathbf{S}\mathbf{v} = \mathbf{P}\mathbf{v}_C \quad \text{with} \quad \mathbf{v} = \begin{pmatrix} \mathbf{v}_F \\ \mathbf{v}_C \end{pmatrix} \quad (10)$$

must hold for an arbitrary vector \mathbf{v} where the variables and equations have been permuted such that F-variables come first. (Those variables which are to be contained in the coarse grid are commonly referred to as C-variables whereas the complementary set is related to the F-variables.)

In this talk we present a formalism for the construction of proper prolongation operators in order to accomplish the transfer from coarse to fine grids within a multigrid algorithm. The main idea is to force relation (10) only locally (i.e. at each coarse variable i) for a certain subset of (algebraically) smooth basis vectors $\mathbf{b}^{k(i)}$ ($k(i) = 1, \dots, n(i)$). Following this idea, classical matrix-dependent prolongations or prolongation operators based on smoothed aggregation can be recovered. Moreover, this rather general framework gives rise to many other variants whose complexity and accuracy can be controlled by a careful selection of basis vectors and smoothing operators, which are applied to construct \mathbf{P} . For example, geometric or other problem-dependent information can be easily exploited to tailor the prolongation to the particular application at hand.

At the end of the talk we will present numerical results for several two- and three-dimensional diffusion problems including jumping coefficients. In particular a real-world three-dimensional image segmentation problem from a medical application is discussed.

Gabriel Wittum
**Modeling and Numerical Simulation of Biological Systems by Multigrid
Methods**

Kettenhofweg 139
60325 Frankfurt am Main
`wittum@g-csc.de`
Gabriel Wittum
G-CSC, University of Frankfurt

A typical feature of biological systems is their high complexity and variability. This makes modelling and computation very difficult, in particular for detailed models based on first principles. The problem starts with modelling geometry, which has to extract the essential features from those highly complex and variable phenotypes and at the same time has to take in to account the stochastic variability. Moreover, models of the highly complex processes which are going on these geometries are far from being well established, since those are highly complex too and often couple on a hierarchy of scales in space and time. Simulating such systems always puts the whole approach to test, including modeling, numerical methods and software implementations. In combination with validation based on experimental data, all components have to be enhanced to reach a reliable solving strategy.

To handle problems of this complexity, new mathematical methods and software tools are required. In recent years, new approaches such as parallel adaptive multigrid methods and corresponding software tools have been developed allowing to treat problems of huge complexity. In the lecture we report on the numerical simulation of the diffusion of xenobiotics through human skin. First computations for this problem were made in the last decade, yielding new insight into permeation pathways through human skin, which were confirmed experimentally ten years later.

Chao Yang
**Parallel Two-level Schwarz Methods for Fully Implicit Solution of
Shallow Water Equations on the Cubed-sphere**

Department of Computer Science
University of Colorado at Boulder
Boulder
CO 80309
`chao.yang@colorado.edu`
Xiao-Chuan Cai

Department of Computer Science, University of Colorado at Boulder, Boulder, CO 80309

In this talk a parallel two-level Schwarz preconditioner is proposed as part of a fully implicit solver for the shallow water equations discretized with a finite volume scheme on the cubed-sphere. With such an implicit time integration scheme the time step size is no longer constrained by the CFL condition, which is usually required when using explicit or semi-implicit methods. The price to pay is that a large sparse nonlinear system of equations has to be solved at every time step. When the one-level method is used, our numerical experiments show that the number of nonlinear iterations per time

step is almost independent of the number of processors and the time step size, but the number of linear iterations grows when we increase the number of processors. In this work we introduce a coarse grid correction to the Schwarz preconditioner, and we show numerically that, with the help of the second level, the number of linear iterations is nearly independent of the number of processors and is less sensitive to the implicit time step size. We present numerical results obtained on machines with thousands of processors.

Ulrike Meier Yang
Improving Interpolation for Aggressive Coarsening

Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
Box 808
L-560
Livermore
CA 94551
`umyang@llnl.gov`

Aggressive coarsening has been a very effective tool to reduce computational complexities as well as memory requirements in algebraic multigrid, however it requires an interpolation operator that deals effectively with the long distances between coarse grid points. Usually aggressive coarsening is combined with multipass interpolation, an interpolation based on direct interpolation, which leads to deteriorating convergence and decreased numerical scalability.

Various new long range interpolation operators based on distance-two interpolation operators, which have a higher of accuracy than direct interpolation, are investigated. This presentation describes the resulting algorithms, their convergence behavior, and analyzes their performance on parallel computers.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-ABS-409876

Irad Yavneh
Cyclic Reduction Multigrid Revisited

Department of Computer Science
Technion - Israel Institute of Technology
Haifa 32000
Israel
`irad@cs.technion.ac.il`
Chen Greif
University of British Columbia, Canada

We revisit an early 90's technique due to Golub and Tuminaro, originally suggested for two-dimensional convection-diffusion. The idea is to perform a single cyclic reduction step before reverting to multigrid. We generalize this idea and introduce a novel smoothing analysis for the resulting problem, which yields explicit analytical results. In particular, we prove for $2d+1$ point "star" stencils, where d is the dimension, that the suitably defined h -ellipticity measure is always increased by the cyclic

reduction step for symmetric M -matrices. Additionally, we introduce a novel relaxation method employing a cyclic reduction step, dubbed cyclic reduction relaxation (CRR), and analyze its smoothing properties, which turn out to be exceptional. Numerical computations show a close agreement between the smoothing analysis results and actual V -cycle convergence factors.