Mathematisches Forschungsinstitut Oberwolfach

Report No. 24/2014

DOI: 10.4171/OWR/2014/24

Schnelle Löser für Partielle Differentialgleichungen

Organised by
Randolph E. Bank, La Jolla
Lars Grasedyck, Aachen
Wolfgang Hackbusch, Leipzig
Gabriel Wittum, Frankfurt

May 11 - May 16, 2014

Abstract. The workshop "Schnelle Löser für partielle Differentialgleichungen", organised by Randolph E. Bank (La Jolla), Lars Grasedyck (Aachen), Wolfgang Hackbusch (Leipzig) and Gabriel Wittum (Frankfurt), was held May 11 - May 16, 2014. This meeting was well attended by 52 participants with broad geographic representation from 11 countries and 3 continents. This workshop was a nice blend of researchers with various backgrounds.

Mathematics Subject Classification (2010): 65N55, 65M55, 65M60, 65N30, 65N50, 65F10, 65F50, 65H10, 65K10.

Introduction by the Organisers

The conference was organized by Randolph E. Bank, UCSD, La Jolla, Lars Grasedyck, RWTH Aachen, Wolfgang Hackbusch, MPI Leipzig, and Gabriel Wittum, University of Frankfurt. This was the seventh one in a series of conferences on fast solvers held at Oberwolfach since 1999. The idea of these workshops is to bring together experts from the different thriving areas of solvers and offer a platform for scientific exchange and progress.

The field of solvers for the algebraic systems arising from the discretization of partial differential equations has developed to a major area of numerical mathematics and scientific computing. Solvers are the essential part of simulation codes for problems from science and technology, in many cases determining the complexity of the whole simulation. By virtue of that, the choice of the solver can decide on the realiability of a simulation and if it can be done at all. Thus, solvers are

a substantial mathematical component of most simulation tools and a major contribution of mathematics to quite a lot of applied disciplines. This has increased the interest in mathematics of colleagues from the applied sciencies over the last two decades substantially.

Major areas of solvers represented at the workshop are: Multigrid methods, H-matrices, domain decomposition methods, and conjugate gradient methods and their scalable parallelization on huge numbers of cores. Often these methods are combined, e.g. conjugate gradient like methods are used as accelerator for multigrid. Besides that, several talks were given on other aspects of solving partial differential equations, such as discretization schemes and the algebraic properties of the resulting stiffness matrices, overall solution strategies, and application areas where solving plays a crucial rôle.

The question of the right solver for critical application problems is still open, but new approaches have been developed in recent years. New light is shed on the solver question by the recent change of paradigm in computer architecture. The modern multicore processors with additional strong GPU and MIC accelerators pose a new and serious challenge for the development of fast solvers. A total of 25 presentations gave a nice overview over the current research, open problems and new developments. Intense discussions provided the opportunity to go into details of novel algorithms and approaches.

In multigrid methods, a lot of research is going in the direction of developing robust methods for special applications. This is a challenging topic requiring mathematical expertise as well as understanding of the model and the application process itself. Another major topic is Algebraic Multigrid. AMG methods are already wide spread in several applied communities. However, a lot of open problems remains and the final algorithm is not yet in sight. Several talks also were related to performance issues of multigrid on certain computer architectures such as super scalar or parallel computers. Multigrid research is thriving more than ever.

Another bunch of talks were about domain decomposition methods. These methods are of particular interest for multiphysics problems and parallelization issues. Several new developments have been reported and discussed, giving interesting future perspectives. Often techiques from domain decomposition analysis can be used to analyze other methods e.g. multigrid. A novel technique useful together with domain decomposition and multigrid, but can also stand on its own, are hierarchical matrices (H-matrices). Here, several talks have shown the impressive level of development these methods already have obtained.

Tensor representation and solving the corresponding equations was discussed in several talks at the conference. This novel class of numerical reduction methods got a lot of attention. It comes from solving high dimensional problems, but can be used to reduce also low dimenisonal ones. Further talks have discussed solver techniques for application problems as well as other problem areas like optimization.

In total, the workshop was very successful in bringing together international-level experts from different areas and disciplines. Meanwhile, the Oberwolfach workshop on "Schnelle Löser für partielle Differentialgleichungen" is established as major event in the solver community and a mainstay for novel developments.

Acknowledgement: The MFO and the workshop organizers would like to thank the National Science Foundation for supporting the participation of junior researchers in the workshop by the grant DMS-1049268, "US Junior Oberwolfach Fellows".

Workshop: Schnelle Löser für Partielle Differentialgleichungen

Table of Contents

Abstracts

Adaptivity and Preconditioning for High-Dimensional Elliptic Partial Differential Equations

Markus Bachmayr
(joint work with Wolfgang Dahmen)

We consider high-dimensional operator equations $Au = f$ with elliptic $A\colon \mathcal{H} \to \mathcal{H}'$, where $\mathcal{H}'$ is a separable Hilbert space. In [2], we have analysed an adaptive method for the construction of low-rank tensor approximations of u based on the hierarchical tensor format [4, 5]. Our analysis required the availability of suitable low-rank representations (or approximations) of A. In this regard, some additional obstructions arise when $\mathcal{H}$ is not endowed with a cross norm, for instance in the case of Sobolev spaces. In this talk, we focus in particular on the treatment of second-order problems. Specifically, we consider the model case

$$A = -\sum_{i,j=1}^{d} a_{ij}\partial_i\partial_j + a_0\mathrm{id}$$

with $\mathcal{H} = H^1_0(\Omega)$, $\Omega = (0,1)^d$; the coefficients a_0, a_{ij} are assumed to be constant, although this is not essential.

The adaptive scheme is based on a Riesz basis representation of the problem. To obtain a Riesz basis in the present case, we may start from a tensor product orthonormal wavelet basis $\{\Psi_\nu\}_{\nu\in\mathbb{N}^d}$ of $L_2(\Omega)$, where $\Psi_\nu = \psi_{\nu_1}\otimes\cdots\otimes\psi_{\nu_d}$. A Riesz basis of $H^1_0(\Omega)$ (assuming sufficient regularity and a suitable boundary adaptation of the ψ_{ν_i}) can then be obtained by rescaling: let

$$\omega_\nu := \|\Psi_\nu\|_{H^1} \sim \Big(\sum_{i=1}^{d} 2^{2j(\nu_i)}\Big)^{\frac{1}{2}},$$

where $j(\nu_i)$ denotes the wavelet level of ν_i, then $\{\omega_\nu^{-1}\Psi_\nu\}_{\mathbb{N}^d}$ is a Riesz basis of $H^1_0(\Omega)$ with, as noted in [3], d-independent condition number. Defining the infinite matrices $\mathbf{S} := \mathrm{diag}(\omega_\nu)$, $\mathbf{T} := (\langle A\Psi_\nu, \Psi_\mu\rangle)_{\mu,\nu}$, the basis representation of the original problem reads $\mathbf{S}^{-1}\mathbf{T}\mathbf{S}^{-1}\mathbf{u} = \mathbf{f}$ with $\mathbf{f}_\nu = \omega_\nu^{-1}\langle f, \Psi_\nu\rangle$ and $\mathbf{u}_\nu = \omega_\nu\langle u, \Psi_\nu\rangle$. Our aim is, for any given $\varepsilon > 0$, to find u_ε such that $\|u - u_\varepsilon\|_{H^1(\Omega)} \le \varepsilon$. In the basis representation, the latter translates to $\|\mathbf{u} - \mathbf{u}_\varepsilon\|_{\ell_2(\mathbb{N}^d)} \le C_\Psi\,\varepsilon$. On the one hand, we need to find a suitable finite set of basis indices $\Lambda_\varepsilon = \Lambda_\varepsilon^{(1)} \times \cdots \times \Lambda_\varepsilon^{(d)} \subset \mathbb{N}^d$ for $\mathbf{u}_\varepsilon$, on the other hand we use a nonlinear representation in hierarchical tensor format for the high-dimensional coefficient vector $(\mathbf{u}_{\varepsilon,\nu})_{\nu\in\Lambda_\varepsilon}$, where we also need to identify appropriate tensor ranks. We construct $\mathbf{u}_\varepsilon$ by an iterative scheme with all operations performed in the hierarchical format, which requires compatible representations of the involved operators $\mathbf{T}$ and $\mathbf{S}^{-1}$.

Whereas $\mathbf{T}$ – for suitably structured (a_{ij}) – has a representation of low hierarchical rank, e.g. of maximum hierarchical rank two in the case $A = -\Delta$, we are facing the problem that on the full index set $\mathbb{N}^d$, the rank of $\mathbf{S}^{-1}$ is unbounded.

With $\mathbf{S}^{-1}$ playing the role of a preconditioner, this can essentially be seen as a consequence of the spectral properties of the Laplacian.

As in [2] (and the precursor in [1]), our adaptive scheme is based on a perturbed Richardson iteration, which now also requires a suitable low-rank approximation of $\mathbf{S}^{-1}$. Since a direct approximation of $\mathbf{S}^{-1}$ leads to unnecessarily high ranks, we replace it by an equivalent diagonal scaling operator $\tilde{\mathbf{S}}^{-1}$, which for a fixed $\delta \in (0,1)$ satisfies $(1-\delta)\|\mathbf{S}^{-1}\mathbf{v}\| \leq \|\tilde{\mathbf{S}}^{-1}\mathbf{v}\| \leq (1+\delta)\|\mathbf{S}^{-1}\mathbf{v}\|$ for any finitely supported $\mathbf{v}$, and which has more efficient rank-n approximations $\tilde{\mathbf{S}}_n^{-1}$. With a tensor recompression operation P_ε and a tensor coarsening operation C_ε as in [2], the method can be written briefly as

$$\mathbf{u}_{m+1} = \mathrm{C}_{\varepsilon_{1,m}} \mathrm{P}_{\varepsilon_{2,m}} \big(\mathbf{u}_j - \omega(\tilde{\mathbf{S}}_{n_m}^{-1} \tilde{\mathbf{T}}_{\varepsilon_{3,m}} \tilde{\mathbf{S}}_{n_m}^{-1} \mathbf{u}_m - \tilde{\mathbf{f}}_{\varepsilon_{4,m}})\big)$$

where $\omega > 0$ is a scaling factor, $\tilde{\mathbf{f}}_\varepsilon$ is an approximation of $\mathbf{f}$, and $\tilde{\mathbf{T}}_\varepsilon$ is an approximation of the infinite matrix $\mathbf{T}$ by wavelet compression of its lower-dimensional tensor components. Here we need to choose the parameters n_m and $\varepsilon_{1,m}, \dots, \varepsilon_{4,m}$ to both ensure convergence to the exact solution and retain control over the complexity of the iteration. We construct $\tilde{\mathbf{S}}^{-1}$ by a suitable exponential sum approximation with relative error bound, based on sinc quadrature, of $t \mapsto t^{-\frac{1}{2}}$ on $[1,\infty)$, which yields approximations

$$\tilde{\mathbf{S}}_n^{-1} = \sum_{k=1}^{n} \mathbf{S}_k^{(1)} \otimes \cdots \otimes \mathbf{S}_k^{(d)}$$

with lower-dimensional diagonal operators $\mathbf{S}_k^{(i)}$ and $\tilde{\mathbf{S}}^{-1} = \lim_{n\to\infty} \tilde{\mathbf{S}}_n^{-1}$. As a result, we arrive at a condition on the required approximation ranks of the form

$$n_j \geq n_0(\delta) + c\big(|\ln \varepsilon_{4,m}| + \max_{\substack{i=1,\dots,d \\ \{\nu\colon \mathbf{u}_{m,\nu}\neq 0\}}} j(\nu_i)\big)$$

with constants n_0, c depending in particular on δ and $\mathbf{T}$. In other words, the maximum active wavelet level in the current iterate $\mathbf{u}_m$ enters linearly in the required approximation rank n_m.

In order to obtain a complexity analysis of the scheme similarly to [2], we thus need to control the maximum wavelet levels that can arise in the course of the iteration. In our present setting, this can be achieved if $f \in H^{-1+t}(\Omega)$ for some $t > 0$; with this additional regularity condition, we subsequently obtain operation complexity estimates along the lines of [2] also in the present setting of problems posed on Sobolev spaces. Under model assumptions that are satisfied, e.g., for certain high-dimensional Poisson problems, the resulting bound on the number of operations required for constructing u_ε with $\|u_\varepsilon - A^{-1}f\|_{H^1} \leq \varepsilon$ is of the form $C_d|\ln\varepsilon|^{K_d}\varepsilon^{-\frac{1}{s}}$. Here s is the approximation order for the one-dimensional tensor components, K_d has an affine linear dependence on $\ln d$, and C_d grows superalgebraically, but subexponentially in d. Numerical experiments with our method for the problem $-\Delta u = 1$ on $\Omega = (0,1)^d$ with homogeneous Dirichlet boundary conditions demonstrate that the number of operations required to guarantee a certain error tolerance in $H^1(\Omega)$ in fact shows a low-degree polynomial scaling in d.

References

[1] M. Bachmayr, *Adaptive Low-Rank Wavelet Methods and Applications to Two-Electron Schrödinger Equations*, PhD thesis, RWTH Aachen (2012).

[2] M. Bachmayr and W. Dahmen, *Adaptive Near-Optimal Rank Tensor Approximation for High-Dimensional Operator Equations*, to appear in Foundations of Computational Mathematics.

[3] T.J. Dijkema, Ch. Schwab, and R. Stevenson, *An adaptive wavelet method for solving high-dimensional elliptic PDEs*, Constructive Approximation **30** (2009), 423–455.

[4] W. Hackbusch, S. Kühn, *A New Scheme for the Tensor Representation*, Journal of Fourier Analysis and Applications **15** (2009), 706–722.

[5] W. Hackbusch. *Tensor Spaces and Numerical Tensor Calculus*, volume 42 of *Springer Series in Computational Mathematics*, Springer-Verlag Berlin Heidelberg (2012).

Multigrid methods for saddle point problems

Susanne C. Brenner

(joint work with Hengguang Li, Duk-Soon Oh, Li-Yeng Sung)

We present a general framework for the design and analysis of multgrid methods for saddle point problems of the form

$$\begin{aligned} a(u,v) + b(v,p) &= F(v) \qquad \forall\, v \in V \\ b(u,q) - c(p,q) &= G(q) \qquad \forall\, v \in Q \end{aligned}$$

that are associated with mixed finite element discretizations of elliptic boundary value problems [1]. These multigrid methods are uniformly convergent in the energy norm on general polyhedral domains Ω.

For the Stokes and Lamé systems discretized by the Taylor-Hood mixed finite element methods, these multgrid methods converge uniformly in the energy norm

$$\|u\|_{H^1(\Omega)} + \|p\|_{L_2(\Omega)}$$

For the Darcy system discretized by the Raviart-Thomas-Nédélec mixed finite element methods of order at least 1, these multigrid methods converge uniformly in the energy norm

$$\|u\|_{L_2(\Omega)} + \|p\|_{H^1(\Omega;\mathcal{T}_h)}$$

where $\|\cdot\|_{H^1(\Omega,\mathcal{T}_h)}$ is a standard DG norm for discontinuous piecewise polynomial functions.

The key ingredients for these multigrid methods are smoothers that involve block diagonal preconditioners. The preconditioner is chosen so that the scale of mesh-dependent norms associated with the smoothers has the following properties.

(1) The smoother enjoys standard smoothing properties with respect to the scale of mesh-dependent norms.

(2) The scale of mesh-dependent norms is equivalent to a scale of Sobolev norms that include the energy norm for the saddle point problem and the fractional order Sobolev norm associated with the regularity of the elliptic boundary value problem.

The second property ensures that approximation properties with respect to the mesh-dependent norms can be established without full elliptic regularity, which

together with the first property yields the uniform convergence of W-cycle algorithms in the energy norm for general polyhedral domains if the number of smoothing steps is sufficiently large. Numerical results indicate that the V-cycle algorithm is also uniform convergent.

For the Stokes and Lamé systems, the preconditioner for the smoothing steps involves an optimal preconditioner for the discrete Laplace operator associated with continuous Lagrange finite element spaces. For the Darcy system it involves an optimal preconditoner for the discrete Laplace operator associated with discontinuous finite element spaces.

The general framework is applicable to problems with nonsymmetric $a(\cdot,\cdot)$ (for example, the Oseen problem) and problems with nonsymmetric $c(\cdot,\cdot)$ (for example, the diffusion-convection/advection-reaction problems).

Other examples include mixed finite element methods for the biharmonic problem using Lagrange elements [2] and mixed finite element methods for linear elasticity in the stress-displacement formulation where the symmetry of the stress is only weakly enforced [3].

Details of this work can be found in [4, 5].

References

[1] D. Boffi, F. Brezzi, and M. Fortin, *Mixed Finite Element Methods and Applications*, Springer, Heidelberg, 2013.

[2] I. Babuška, J. Osborn, and J. Pitkäranta, *Analysis of mixed methods using mesh dependent norms*, Math. Comp.,**35** (1980), 1039–1062.

[3] J. Guzmán, *A unified analysis of several mixed methods for elasticity with weak stress symmetry* J. Sci. Comput., **44** (2010), 156–169.

[4] S.C. Brenner, H. Li, and L.-Y. Sung, *Multigrid methods for saddle point problems*: *Stokes and Lamé systems*, Numer. Math., published online January 2014. (DOI: 10.1007/s00211-014-0607-3).

[5] S.C. Brenner, D.-S. Oh and L.-Y. Sung, *Multigrid methods for saddle point problems*: *Darcy system*, preprint, 2014.

$\mathcal{H}^2$-matrix preconditioners

Steffen Börm

(joint work with Knut Reimer)

Hierarchical matrices [13, 15, 11, 14] can be used to construct preconditioners for FEM and BEM problems: generalized regularity results [2, 5, 10] imply that the solution operators of strongly elliptic PDEs can locally be approximated by low-rank operators. Storing these operators in factorized form gives rise to the representation as a hierarchical matrix.

Of particular interest are approximations of triangular factorizations like the LR or Cholesky decomposition, since these lead to reduced storage requirements and setup times compared to the approximation of the entire inverse. Experiments show that this approach leads to efficient and very robust preconditioners [17, 1,

12], and this observation has recently be confirmed by theoretical investigations [10, 9].

$\mathcal{H}^2$-matrices [16, 7, 6] are a variant of hierarchical matrices that employ an additional compression to reduce storage requirements and compute times even further. It has been proven [4] that standard approximation results for hierarchical matrices can easily be extended to $\mathcal{H}^2$-matrices, so we are faced with the question of whether we can construct $\mathcal{H}^2$-matrix approximations by algorithms that are similar to those used for hierarchical matrices.

This talk, based on [8], shows that we can approximate products, inverses and factorizations of $\mathcal{H}^2$-matrices using only two fundamental operations:

- the simultaneous multiplication of an $\mathcal{H}^2$-matrix with k vectors and
- the (approximate) update of an $\mathcal{H}^2$-matrix by a factorized rank-k-matrix.

This approach leads to algorithms that a structurally similar to their counterparts for standard hierarchical matrices, but are significantly more efficient: the fundamental operations can be performed in linear complexity and give rise to algorithms of total complexity $\mathcal{O}(nk^2 \log n)$ for approximating the inverse or the triangular factorization of an $\mathcal{H}^2$-matrix.

Numerical experiments illustrate that this approach leads to preconditioners for FEM and BEM problems that require $\mathcal{O}(n)$ units of storage, and $\mathcal{O}(n \log n)$ operations to set up.

Besides the construction of preconditioners, approximate $\mathcal{H}^2$-matrix factorizations can also be employed to compute eigenvalues of elliptic partial differential operators by "slicing the spectrum" [18]. Numerical experiments [3] show that one step of the resulting method takes only $\mathcal{O}(n \log n)$ operations.

The last part of the talk is devoted to a possible future application of the method: if the cluster tree is chosen correctly, we should be able to compute approximate factorizations of the stiffness matrices of saddle-point problems arising from the mixed finite element discretization of Darcy's equation.

References

[1] M. Bebendorf. Why finite element discretizations can be factored by triangular hierarchical matrices. *SIAM J. of Numer. Anal.*, 45(4):1472–1494, 2007.

[2] M. Bebendorf and W. Hackbusch. Existence of $\mathcal{H}$-matrix approximants to the inverse FE-matrix of elliptic operators with L^∞-coefficients. *Numer. Math.*, 95:1–28, 2003.

[3] P. Benner, S. Börm, T. Mach and K. Reimer. Computing the eigenvalues of symmetric $\mathcal{H}^2$-matrices by slicing the spectrum. Preprint MPIMD/14-06, Max-Planck-Institut für Dynamik komplexer technischer Systeme, 2014. available at `http://arxiv.org/abs/1403.4142`.

[4] S. Börm. Data-sparse approximation of non-local operators by $\mathcal{H}^2$-matrices. *Lin. Alg. Appl.*, 422:380–403, 2007.

[5] S. Börm. Approximation of solution operators of elliptic partial differential equations by $\mathcal{H}$- and $\mathcal{H}^2$-matrices. *Numer. Math.*, 115(2):165–193, 2010.

[6] S. Börm. *Efficient Numerical Methods for Non-local Operators: $\mathcal{H}^2$-Matrix Compression, Algorithms and Analysis*, volume 14 of *EMS Tracts in Mathematics*. EMS, 2010.

[7] S. Börm and W. Hackbusch. Data-sparse approximation by adaptive $\mathcal{H}^2$-matrices. *Computing*, 69:1–35, 2002.

[8] S. Börm and K. Reimer. Efficient arithmetic operations for rank-structured matrices based on hierarchical low-rank updates. Technical Report, available at http://arxiv.org/abs/1402.5056.
[9] M. Faustmann, J. M. Melenk, and D. Praetorius. Existence of $\mathcal{H}$-matrix approximants to the inverses of BEM matrices: the simple-layer operator. Technical Report 37, Institut für Analysis und Scientific Computing, TU Wien, 2013. available at http://www.asc.tuwien.ac.at/preprint/2013/asc37x2013.pdf.
[10] M. Faustmann, J. M. Melenk, and D. Praetorius. $\mathcal{H}$-matrix approximability of the inverse of FEM matrices. Technical Report 20, Institut für Analysis und Scientific Computing, TU Wien, 2013. available at http://www.asc.tuwien.ac.at/preprint/2013/asc20x2013.pdf.
[11] L. Grasedyck and W. Hackbusch. Construction and arithmetics of $\mathcal{H}$-matrices. *Computing*, 70:295–334, 2003.
[12] L. Grasedyck, R. Kriemann, and S. LeBorne. Domain decomposition based $\mathcal{H}$-LU preconditioning. *Numer. Math.*, 112(4):565–600, 2009.
[13] W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices. *Computing*, 62:89–108, 1999.
[14] W. Hackbusch. *Hierarchische Matrizen — Algorithmen und Analysis*. Springer, 2009.
[15] W. Hackbusch and B. N. Khoromskij. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part II: Application to multi-dimensional problems. *Computing*, 64:21–47, 2000.
[16] W. Hackbusch, B. N. Khoromskij, and S. A. Sauter. On $\mathcal{H}^2$-matrices. In H. Bungartz, R. Hoppe, and C. Zenger, editors, *Lectures on Applied Mathematics*, pages 9–29. Springer-Verlag, Berlin, 2000.
[17] M. Lintner. The eigenvalue problem for the 2d Laplacian in $\mathcal{H}$-matrix arithmetic and application to the heat and wave equation. *Computing*, 72:293–323, 2004.
[18] B. Parlett. The Symmetric Eigenvalue Problem. Prentice-Hall, Englewood Cliffs (1980)

Affine conjugate adaptive Newton methods for nonlinear mechanics

Peter Deuflhard
(joint work with Martin Weiser)

The talk presented an adaptive Newton method in function space for nonlinear mechanics. The algorithm is run in three levels:

(I) outer iteration: global Newton method in function space with adaptive damping, where the arising cubic upper bound is computationally estimated, thus leading to a predictor and a corrector strategy;
(II) adaptive multilevel discretization of function minimization, both convex and polyconvex, where the discretization errors on each level are controlled by a-posteriori estimates; the energy error norm is estimated in terms of the (convex) functional for linear mechanics;
(III) inner iteration: preconditioned conjugate gradient method (PCG), applied to the indefinite case of nonlinear mechanics by either truncated or regularized PCG.

All three levels are intertwined, e.g., the mesh generation takes the nonlinearity into account. This algorithm has been successfully applied in challenging problems of cranio-maxillo facial surgery planning and should work in the general context of nonlinear mechanics.

References

[1] P. Deuflhard: *Newton Methods for Nonlinear Problems. Affine Invariance and Adaptive Algorithms*, Springer Series in Computational Mathematics **35** (2006).

[2] P. Deuflhard, M. Weiser: *Local inexact Newton multilevel FEM for nonlinear elliptic problems*. In M.-O. Bristeau, G. Etgen, W. Fitzgibbon, J.-L. Lions, J. Periaux, and M. Wheeler, editors. *Computational Science for the 21st Century*, Wiley-Interscience-Europe (1997), 129–138.

[3] P. Deuflhard, M. Weiser: *Global inexact Newton multilevel FEM for nonlinear elliptic problems*. In W. Hackbusch and G. Wittum, editors. *Multigrid Methods*, *Lecture Notes in Computational Science and Engineering* **3**, Springer (1998), 71–89.

[4] M. Weiser, P. Deuflhard, B. Erdmann: *Affine conjugate adaptive Newton methods for nonlinear elastomechanics*. OMS **22** (2007), 413–431.

The Laplace Transform versus Parareal

Craig C. Douglas

(joint work with Li Deng, Hyoseop Lee, Imbunm Kim, and Dongwoo Sheen)

Both the Laplace Transform and the Parareal family of algorithms promise to provide completely parallel in time and space computational results. Given a random problem it is unclear which algorithm will run faster. In this talk, we provide the following:

- Define the algorithms in question, including more specific details than usual.
- Define interesting parallel environments, some of which do not exist yet.
- Demonstrate several computational environments in which one of the algorithms can be expected to be faster than the other algorithm, including when compact 4th and 6th order finite difference schemes are used.

Details can be found in [1] and [2].

References

[1] C. C. Douglas, L. Deng, H. Lee, and D. Sheen, *A comparison of the Parareal and the Laplace transform algorithms*, in Proceedings of DCABES 2012, Guilin, China, 2012, Q. Guo and C. C. Douglas (eds.), IEEE Computer Society CPS, Los Alamitos, CA, 2012, pp. 1–4.

[2] C. C. Douglas, H. Lee, I. Kim, and D. Sheen, *Higher-order schemes for the Laplace transformation method for parabolic problems*, Computation and Visualization in Science, 14 (2011), pp. 14–37.

Parallel Multigrid in Space and Time

Robert D. Falgout
(joint work with S. Friedhoff, Tz. V. Kolev, S. P. MacLachlan, J. B. Schroder)

Multigrid methods are important techniques for efficiently solving huge linear systems and they have already been shown to scale effectively on millions of cores. Future exascale architectures will require solvers to exhibit extreme levels of concurrency (1B cores), minimize data movement, exploit machine heterogeneity, and demonstrate resilience to faults. Multigrid approaches are ideal for addressing these challenges. In this talk, we discuss two areas of multigrid research. The first is a new non-Galerkin approach for constructing coarse-grid operators in algebraic multigrid (AMG) [1]. The method uses a stencil-collapsing technique to reduce the number of nonzeros in the Galerkin matrix without destroying convergence. This greatly improves parallel performance by reducing communication on coarser grids. The second research topic is a parallel time integration approach based on multigrid reduction (MGR) techniques [2]. The advantage of the approach is that it is easily integrated into existing codes because it simply calls the users existing time-stepping routine. Parallel results show the potential for significant speedup over standard time stepping for parabolic equations.

References

[1] Robert D. Falgout and Jacob B. Schroder, Non-galerkin coarse grids for algebraic multigrid, *SIAM J. Sci. Comput.*, to appear. LLNL-JRNL-641635.
[2] R. D. Falgout, S. Friedhoff, Tz. V. Kolev, S. P. MacLachlan, and J. B. Schroder, Parallel time integration with multigrid, *SIAM J. Sci. Comput.*, submitted. LLNL-JRNL-645325.

Models and Numerical Strategies for Inelastic Processes in Industrial and Biological Problems

Alfio Grillo
(joint work with Raphael Prohl, Gabriel Wittum)

The theory of inelastic processes constitutes an intensively investigated research field of theoretical and computational Mechanics. It addresses the characterisation of the mechanical behaviour of continuum bodies that, besides changing shape, may also undergo inelastic reorganisations of their internal structure in response to stimuli of various nature and acting at different length and time scales.

The study of the structural evolution of continuum bodies covers different classes of phenomena, these ranging from the onset of plastic deformations in the building materials used in industrial applications to the growth and remodelling of biological tissues. The physics lying behind these phenomena is very different, since it depends on the material properties of the considered body and on the type of interactions that the body itself experiences with its environment. Nevertheless, the mathematical models and the computational strategies, which are elaborated to formulate and solve inelastic problems, have many features in common.

The scope of this contribution is to review and extend some already existing models of inelastic phenomena, and to propose a computational procedure capable of solving accurately problems arising both in the industrial context and in the biomechanical one. The proposed algorithm, which should serve as a basis for an efficient structural mechanical solver, has been tested by solving some benchmarks available in literature and largely adopted in finite-strain elastoplasticity. For this reason, and for the sake of conciseness, the methodologies discussed in the following shall be dealing only with the non-linear elastoplastic modelling of metals.

The theoretical and computational tools adopted to approach inelastic processes are often inspired by the Theory of Finite Elastoplasticity (cf., e.g., [1]), where the deformation gradient $\boldsymbol{F}$, which accounts for the overall change of shape of the studied body (cf. [2] for details), is decomposed multiplicatively into an elastic and a plastic part, i.e., $\boldsymbol{F} = \boldsymbol{F}_\mathrm{e}\boldsymbol{F}_\mathrm{p}$. This decomposition has its roots in the micro-mechanical description of plasticity. For example, in the case of single-crystal plasticity, $\boldsymbol{F}_\mathrm{e}$ accounts for the deformation of the crystal lattice, whereas $\boldsymbol{F}_\mathrm{p}$ is said to move "the material through the lattice via dislocation motion" [3]. In Biomechanics, the tensor $\boldsymbol{F}_\mathrm{p}$ describes the inelastic deformations associated with the material inhomogeneities that accompany the remodelling of a tissue [4], that is its structural reorganisation in response to applied loads (the latter one tends to optimise the stress distribution inside the tissue).

The plastic part of the overall deformation, $\boldsymbol{F}_\mathrm{p}$, can be viewed as an independent tensor variable. Therefore, in a purely mechanical framework, the independent kinematic descriptors of a body undergoing an elastoplastic process are $\boldsymbol{F}$ and $\boldsymbol{F}_\mathrm{p}$, while the elastic part of the overall deformation is given by $\boldsymbol{F}_\mathrm{e} = \boldsymbol{F}\boldsymbol{F}_\mathrm{p}^{-1}$. This interpretation of $\boldsymbol{F}_\mathrm{p}$ follows fundamental concepts presented in [5, 6], where $\boldsymbol{F}_\mathrm{p}$ acquires the meaning of the Lagrange parameter associated with the body's plastic degrees of freedom.

Many elastoplastic models introduce, along with $\boldsymbol{F}_\mathrm{p}$, the hardening parameter. In general, this is a second-order tensor that describes the evolution of the elastic range of a given material.

Looking at the literature, the most relevant differences among the available models of Elastoplasticity concern the determination of the *flow rule*, which is the evolution law for $\boldsymbol{F}_\mathrm{p}$, and the description of the hardening. For example, many models of metal plasticity are based on associative flow rules (i.e., the direction of the plastic strain rate is perpendicular to the yield surface [7]), while, in other circumstances, as is the case for some problems in soil mechanics, plastic flow needs to be modelled by non-associative flow rules. Moreover, depending on the problem that has to be solved, the hardening can be isotropic, kinematic, or a combination of both.

A quite commonly adopted procedure to solve plasticity models is the Return Mapping Algorithm (RMA). The RMA is classically formulated as a closest point projection method. In [3], it is presented for the case of associative and rate-independent plastic flow rules, and under the assumption of isotropic elastoplastic material behaviour. This latter hypothesis allows to formulate the elastoplastic

problem in terms of the symmetric and positive-definite tensor $\boldsymbol{B}_\mathrm{p} = (\boldsymbol{F}_\mathrm{p}^\mathrm{T}.\boldsymbol{F}_\mathrm{p})^{-1}$, which leads to important computational simplifications.

Some modifications of the RMA have been proposed (cf., e.g., [8]) to include rate-dependent elastoplastic behaviour and non-associative flow rules. Moreover, other algorithms, developed more recently, are elaborations of general procedures known as Active Set, Interior Point, and Sequential Quadratic Programming (SQP) methods [9, 10], and rely on the possibility of reformulating the elastoplastic problem as a constrained optimisation problem. For instance, the SQP method developed by Wieners [9] iterates along a sequence of linearised quadratic sub-problems, which consist of the linearised flow rule and the linearised Karush-Kuhn-Tucker (KKT) conditions. The linearisation is performed with respect to the stresses, the plastic multiplier and the displacements, so that only a linearised material response is inserted into the equilibrium equations in every SQP step.

The principal steps of the RMA can be summarised as follows: Firstly, the flow rule and the KKT-conditions are solved independently for every integration point, resulting in a non-linear stress response. Therein, the plastic multiplier, which is necessary to fulfil the KKT-conditions, is computed by a "predictor-corrector" method. Secondly, this response is inserted into the equilibrium equation, which leads to a non-linear variational problem for the displacements. A relevant feature of the RMA is that the weak form of the equilibrium equation is solved together with the flow rule, which is local and, thus, solved only at the integration points.

The RMA, as presented in [3], relies on the concept of *trial stress*, which is obtained by ideally "freezing" the plastic deformation tensor determined at the time step t_n, and setting it equal to the *trial* plastic deformation tensor at time t_{n+1}, i.e., $\boldsymbol{F}_\mathrm{p}^\mathrm{trial}(t_{n+1}) \equiv \boldsymbol{F}_\mathrm{p}(t_n)$. The method introduces a simplification of the flow rule, which facilitates the projection of the trial stress onto the yield surface and, above all, allows to express the unknown $\boldsymbol{B}_\mathrm{p}(t_{n+1})$ entirely in terms of trial quantities. Furthermore, the method is developed for a particularly simple strain energy density function, which is defined as the sum of a term depending solely on $J_\mathrm{e} := \det(\boldsymbol{F}_\mathrm{e})$ (i.e., the volumetric part of $\boldsymbol{F}_\mathrm{e}$) and a term depending solely on $\overline{\boldsymbol{F}}_\mathrm{e} := J_\mathrm{e}^{-1/3}\boldsymbol{F}_\mathrm{e}$ (i.e., the distortional part of $\boldsymbol{F}_\mathrm{e}$).

In this contribution, a computational procedure for elastoplastic problems at finite strains is proposed, which aims to further modify and extend the RMA to much more general flow rules and strain energy density functions. Geometric and kinematic non-linearities as well as non-linear material behaviour are accounted for. The procedure is based on the KKT-system that characterises the elastoplastic material response. This set of equations consists of the equilibrium equation, the flow rule, the (isotropic) hardening law, the KKT-conditions, and the constitutive laws. For ease of exposition, however, hardening is not considered here, and the KKT-multiplier is written in such a form that it can be computed *a posteriori*, once $\boldsymbol{F}$ and $\boldsymbol{B}_\mathrm{p}$ are known. In order to briefly sketch how the algorithm works, no external forces are considered, so that the weak form of the equilibrium equation

and the flow rule are written as follows:

$$\mathcal{W}(\chi, \boldsymbol{B}_\mathrm{p}) = 0, \quad \mathcal{W}(\chi, \boldsymbol{B}_\mathrm{p}) = \int_\Omega \hat{\boldsymbol{P}}(\chi, \boldsymbol{B}_\mathrm{p}) : \mathrm{Grad}(\boldsymbol{w}), \tag{1}$$

$$\mathfrak{G}(\chi, \boldsymbol{B}_\mathrm{p}) = 0, \tag{2}$$

where χ is the deformation (whose "gradient" is $\boldsymbol{F}$), $\mathcal{W}(\chi, \boldsymbol{B}_\mathrm{p})$ is the virtual internal power, $\boldsymbol{w}$ is the virtual velocity, $\hat{\boldsymbol{P}}(\chi, \boldsymbol{B}_\mathrm{p})$ is the first Piola-Kirchhoff stress tensor, and $\mathfrak{G}(\chi, \boldsymbol{B}_\mathrm{p})$ is a compact way of writing the flow rule. Equations (1) and (2) are understood here as already written in time-discrete form.

Given an initial pair $(\chi_0, \boldsymbol{B}_{\mathrm{p}0})$, the proposed procedure linearises both $\mathcal{W}$ and $\mathfrak{G}$ with respect to $\boldsymbol{B}_\mathrm{p}$ in a neighbourhood of $\boldsymbol{B}_{\mathrm{p}0}$ along some plastic increment $\boldsymbol{\Phi}_\mathrm{p}$. This leads to approximated expressions of $\mathcal{W}(\chi, \boldsymbol{B}_\mathrm{p})$ and $\mathfrak{G}(\chi, \boldsymbol{B}_\mathrm{p})$, denoted by $\Delta_\mathcal{W}(\chi, \boldsymbol{B}_{\mathrm{p}0}, \boldsymbol{\Phi}_\mathrm{p})$ and $\boldsymbol{\Delta}_G(\chi, \boldsymbol{B}_{\mathrm{p}0}, \boldsymbol{\Phi}_\mathrm{p})$, respectively, which are linear in $\boldsymbol{\Phi}_\mathrm{p}$. Setting $\boldsymbol{\Delta}_G(\chi, \boldsymbol{B}_{\mathrm{p}0}, \boldsymbol{\Phi}_\mathrm{p})$ equal to zero allows to express $\boldsymbol{\Phi}_\mathrm{p}$ as a function of χ and $\boldsymbol{B}_{\mathrm{p}0}$, i.e., $\boldsymbol{\Phi}_\mathrm{p} = \hat{\boldsymbol{\Phi}}_\mathrm{p}(\chi, \boldsymbol{B}_{\mathrm{p}0})$. Substituting this result into $\Delta_\mathcal{W}$, one obtains the transformed functional $\hat{\Delta}_\mathcal{W}(\chi, \boldsymbol{B}_{\mathrm{p}0})$, which is highly non-linear in χ. Next, $\hat{\Delta}_\mathcal{W}(\chi, \boldsymbol{B}_{\mathrm{p}0})$ is set equal to zero, and the deformation satisfying this equation is determined with the aid of standard iterative methods (e.g., Newton procedure) by linearising in a neighbourhood of χ_0. Then, the whole procedure is repeated until (1) and (2) are satisfied within a prescribed tolerance.

In order to check whether the proposed method, termed *Generalised Plasticity Algorithm* (GPA), delivers results comparable to those obtained by using the RMA method, two benchmarks have been simulated. The first one is a shear-compression test of a unit cube [11] (see fig. 1); the second one is the necking of a cylindrical bar [3] (see fig. 2).

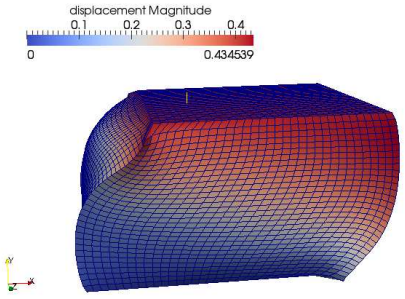

FIGURE 1. Shear compression test of a unit cube at $t = 300\,\mathrm{s}$.

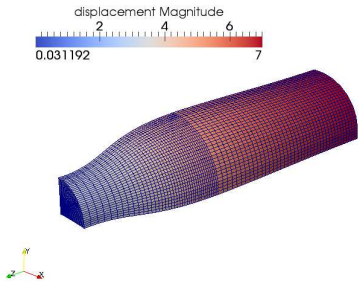

FIGURE 2. Necking of the circular bar in a tensile test. Boundary displacement: 7.0 mm.

Compared to the classical RMA, the proposed computational strategy presents an additional linearisation iteration with respect to the plastic deformations. This implies that the computational effort required by the GPA is bigger than that needed by the classical RMA. Nevertheless, likewise to the RMA, the GPA can be combined with very efficient and robust multigrid methods, and could thus be suited for large problems. The numerical methods used for simulations have been implemented in UG4, a novel version of the software framework UG ("Unstructured Grids") [12]. This toolbox provides fast, massive-parallel solvers for coupled partial differential equations like, e.g., geometric and algebraic multigrid methods. Its new tools for parallel communication (PCL) allow for an efficient scaling of the code on large numbers of processors [13].

References

[1] M.V. Mićunović, *Thermomechanics of Viscoplasticity: Fundamentals and Applications*, Springer-Verlag, Heidelberg (2009).

[2] J.E. Marsden, T.J.R. Hughes, *Mathematical Foundations of Elasticity*, Dove Publications Inc., New York (1983).

[3] J.C. Simo, T.J.R. Hughes, *Computational Inelasticity*, Springer-Verlag, Heidelberg (1998), p. 302.

[4] M. Epstein, G.A. Maugin *Thermomechanics of volumetric growth in uniform bodies*, Int. J. Plast. **16** (2000), 951–978.

[5] P. Cermelli, E. Fried, S. Sellers. *Configurational stress, yield and flow in rate-independent plasticity*, Proc. R. Soc. A., A457 (2001), 1447-1467.

[6] R.A. Eve, B.D. Reddy. *The variational formulation and solution of problems of finite-strain elastoplasticity based on the use of the dissipation function*, Int. J. Numer. Meth. Engng., 37 (1994), 1673-1695.

[7] J. Lubliner. *Normality rules in large-deformation plasticity*, Mech. Mater., 5 (1986), 29-34.

[8] G. Hofstetter, R.L. Taylor. *Non-associative Drucker-Prager plasticity at finite strains*. Communications in Applied Numerical Methods, 6 (1990), 583-589.

[9] C Wieners. *Nonlinear solution methods for infinitesimal perfect plasticity*. Z. angew. Math. Mech., 87 (2007), 643-660.

[10] M. Sauter. *Numerical analysis of algorithms for infinitesimal associated and non-associated elasto-plasticity*. Ph.D. Thesis, Institut für Angewandte und Numerische Mathematik, KIT, Karlsruhe, 2010.

[11] P. Neff, C Wieners. *Comparison of models for finite plasticity: A numerical study*. Comput. Visual. Sci., 6 (2003), 23-25.

[12] A. Vogel, S. Reiter, M. Rupp, A. Nägel, G. Wittum. *UG 4 - A novel flexible software system for simulating PDE based models on high performance computers*. Comput. Visual. Sci., to appear.

[13] S. Reiter, A. Vogel, I. Heppner, M. Rupp, G. Wittum. *A massively parallel geometric multigrid solver in hierarchically distributed grids*. Comput. Visual. Sci., to appear.

GPU accelerated non-linear solvers in biomedical applications

Gundolf Haase
(joint work with Caroline Costa, Manfred Liebmann, Aurel Neic, Gernot Plank)

The presentation is embedded in the context of coupled cardiac electro-mechanical simulations with special focus on the acceleration of non-linear mechanics on accelerator devices [3]. The various solver parts in CARP as elliptic PDE, parabolic PDE, ODEs, or linear elasticity are accelerated on GPUs by using CUDA and it turns out that one GPU substitues 20 CPU cores on a Tier-0 cluster.

The coupling of the bidomain equations with elasticity is fully realized in the CARP code as well as in the Parallel Toolbox (PT). The interfaces for combining CARP and PT have been redesigned regarding flexibility and efficiency which allows solver strategies adapted to non-linear problems. The finite element subroutines for non-linear elasticity are based on the work of C. Augustin [4, 1] and the finite element matrix generation has been partially accelerated by OpenMP and GPU parallelization [2].

Our AMG preconditioner has been extended to elasticity problems with special focus on the non-linear problems. Sophisticated update strategies of the AMG components have been implemented [2] which take into account that the mesh topology will not change during the calculations. This allows to separate matrix graph generation algorithms in finite element generation and AMG operator setup, which are hard to parallelize, from intrinsic parallel parts as recalculation of stiffness matrix entries, intergrid transfer operator entries, coarse operator entries in the non-linear case. The matrix graphs and operator graphs will be generated only once on the CPU while all other parallel recalculations of matrices, coarse operators and interpolation operators can be performed on the accelerator devices.

We did extensive investigations of the AMG solver components for solving the coupled systems resulting from non-linear elasticity. It turned out that the approach for linear elasticity by Griebel/Oeltz/Schweitzer cannot be used without modifications for the linearized elasticity system, especially when the minimal run time is the main criterion. For solving the linear elasticity equations, we achieved the fastest run time, on CPU as well as on GPU, with an unsmoothed block aggregation combined with a block smoother, i.e., the degrees of freedom per node have to be handled together in the smoother but matrix dependent interpolation and restriction is not necessary [1].

The matrix calculation in CARP consumes up to 50% of the compute time. Therefore, further acceleration for electro-mechanics on the GPU can only be achieved if also the matrix calculation in the non-linear iteration is performed completely on the GPU avoiding data transfer between CPU memory and GPU memory. This requires larger changes in CARP and still ongoing work. By defining an element matrix interface between CARP and PT we managed already a very fast stiffness matrix assembling process on the GPU. We did further investigations with finite element matrix calculations for tetrahedral finite elements by using CUDA on the GPU. Vectorizing the code accelerated the element matrix calculations on

one CPU core already by a factor of 5 and the GPU contributed with a further acceleration of 80 [2].

We performed experiments with more general parallelization approaches as OpenACC and OpenMP 4.0 on GPUs and on Intels's Xeon Phi. It turns out that this pragma driven parallelization for many-core environments is getting quite tricky whenever one moves beyond the textbook examples. Nevertheless, there is a high potential for using various accelerator cards with the same code. Currently, the matrix calculation and assembling for tetrahedral elements and the potential problem is used as a test bed for determining the best data structures and algorithms on CPUs, GPUs as well as on the Xeon Phi. The first results are promising taking into account that only one code has to be maintained for all accelerator devices.

References

[1] C. M. Augustin, A. Neic, M. Liebmann, G. Haase, A. Prassl, and G. Plank. Anatomically accurate high resolution modeling of cardiac electromechanics. *Biomechanics and Modeling in Mechanobiology*, under review.

[2] C. Costa, G. Haase, M. Liebmann, A. Neic, and G. Plank. Stepping into fully GPU accelerated biomedical applications. In I. Lirkov, S. Margenov, and J. Wasniewski, editors, *Large Scale Scientific Computing LSSC'13*, volume 8353 of *Lecture Notes in Computer Science*, pages 3–14. Springer, 2014.

[3] A. Neic, M. Liebmann, E. Hoetzl, L. Mitchell, E. J. Vigmond, G. Haase, and G. Plank. Accelerating cardiac bidomain simulations using graphics processing units. *IEEE Trans Biomed Eng*, 59(8):2281–2290, Aug 2012.

[4] G. Plank, A. J. Prassl, and C. M. Augustin. Computational challenges in building multi-scale and multi-physics models of cardiac electro-mechanics. *Biomed Tech (Berl)*, Sep 2013.

Visual Programming and Simulation Workflows

Michael Hoffer
(joint work with Gabriel Wittum)

Historically the application development process consists of three individual steps: editing, compiling and running the application. In most cases the abstract syntax tree and other information generated by the compiler is ignored after compiling and optimizing the application.

However, this information is highly valuable for automatically creating an interactive and visual representation of the application worklfow [1].

Even more important than interactivity is the possibility to create language profiles. This allows for applying mathematical precision to the development process by reducing the language to the minimal subset that is necessary to formulate the desired functionality.

Working with reduced and problem specific language profiles enables the compiler to gain domain knowledge which can be used to automatically eliminate several classes of errors.

We strongly recommend to model the simulation tools and programming languages to handle the increasing complexity of the simulation tools, algorithms and problems users try to solve.

References

[1] M. Hoffer, C. Poliwoda, G. Wittum, *Visual Reflection Library - A Framework for Declarative GUI Programming on the Java Platform*, Computing and Visualization in Science, in press.

Recent advances in grid-based tensor numerical methods in electronic and molecular structure simulations

Boris N. Khoromskij

(joint work with Venera Khoromskaia, Sergey Dolgov)

Grid-based tensor methods provide the efficient tools for numerical approximation of d-dimensional PDEs (discretized on large $n^{\otimes d}$-grids) with linear complexity scaling in the dimension, $O(dn)$. Traditional methods of separable approximation combine the canonical, Tucker, as well as the matrix product state type (MPS) formats and, in particular, TT/HT representations, [1]. The quantized-TT (QTT) approximation [3] is proven to provide the logarithmic data-compression on a wide class of functions and operators. In combinations with the canonical/Tucker/TT formats, it makes possible to efficiently solve multi-dimensional steady-state and dynamical problems by their reformulation in quantized tensor spaces, now with the logarithmic complexity scaling in the size full data on the grid, $O(d \log n)$.

In the proposed talk, we are going to discuss how the grid-based tensor approximation applies to hard problems arising in electronic structure calculations, such as calculation of many-electron integrals, solution of the nonlinear eigenvalue problem for the Hartree-Fock equation, as well as to the equations which model huge lattice-structured compounds, arising in numerical modeling of metals, crystals and polymers [2, 4, 5, 6]. We will also discuss the efficient simultaneous times-space tensor approximation to chemical master equations arising in stochastic modeling of multi-particle reaction processes [7].

Numerical tests indicate the efficiency of newly developed tensor numerical methods on realistic examples in electronic and molecular structure simulation.

http://personal-homepages.mis.mpg.de/bokh

References

[1] B.N. Khoromskij. *Tensors-structured Numerical Methods in Scientific Computing: Survey on Recent Advances.* Chemometr. Intell. Lab. Syst. 110 (2012), 1-19.

[2] B.N. Khoromskij, V. Khoromskaia, and H.-J. Flad. *Numerical solution of the Hartree-Fock equation in multilevel tensor-structured format.* SIAM J. Sci. Comp., **33**(1), 2011, 45-65.

[3] B.N. Khoromskij. *$O(d \log N)$-Quantics Approximation of N-d Tensors in High-Dimensional Numerical Modeling.* J. Constr. Approx. v. 34(2), 257-289 (2011).

[4] V. Khoromskaia, B.N. Khoromskij, and R. Schneider. *Tensor-structured calculation of two-electron integrals in a general basis.* SIAM J. Sci. Comput., **35**(2), 2013, A987-A1010.

[5] V. Khoromskaia. *Black-box Hartree-Fock solver by tensor numerical methods.* Comp. Meth. in Applied Math., vol. 14(2014) No. 1, pp.89-111.
[6] V. Khoromskaia and B.N. Khoromskij. *Grid-based Ewald-type lattice summation of electrostatic potentials by low-rank tensor approximation.* Preprint 116/2013, MPI MiS, Leipzig 2013 (CPC, 2014 to appear).
[7] S. Dolgov, and B.N. Khoromskij. *Simultaneous state-time approximation of the chemical master equation using tensor product formats.* arXiv:1311.3143, 2013; http://arxiv.org/abs/1311.3143, (NLAA, 2014 to appear).

Numerical Solution of Multicomponent Phase Field Systems

Ralf Kornhuber
(joint work with Carsten Gräser, Uli Sack)

Vector-valued phase field models play a prominent role in many practical applications such as, e.g., dynamic recrystallization in solids (non-conserved order parameter) or phase separation of multicomponent alloys (conserved order parameter). We consider phase field models associated with the Ginzburg-Landau free energy

$$\mathcal{E}(u) = \int_\Omega \frac{\varepsilon}{2} \sum_{i=1}^N |\nabla u_i|^2 + \frac{1}{\varepsilon} \Psi_\theta(u)\, dx, \quad \varepsilon > 0, \tag{1}$$

of a vector-valued order parameter $u = (u_1, \ldots, u_N) \in \mathbb{R}^N$ defined on a Lipschitz domain $\Omega \subset \mathbb{R}^d$, $d = 1, 2, 3$. Assuming that the components u_i describe certain concentrations, u is pointwise contained in the closed, convex Gibbs simplex G,

$$u(x,t) \in G = \{v \in \mathbb{R}^N \mid 0 \le v_i,\ v \cdot \mathbf{1} = 1\} \subset \mathbb{R}^N, \quad \mathbf{1} = (1, \ldots, 1)^T. \tag{2}$$

The multi–well potential $\Psi_\theta(u) = \Phi_\theta(u) + \chi_{H_1}(u) + \frac{1}{2}\, Ku \cdot u$ can be decomposed into a convex part $\Phi_\theta + \chi_{H_1}$ and a remaining quadratic part $\frac{1}{2}\, Ku \cdot u$. The convex part consists of the characteristic functional χ_{H_1} enforcing the algebraic constraints $u \in H_1 = \{v \in \mathbb{R}^N \mid v \cdot \mathbf{1} = 1\}$ and of the logarithmic function

$$\Phi_\theta(u) = \sum_{i=1}^N \theta u_i \ln(u_i) \text{ for } \theta > 0, \quad \Phi_\theta(u) = \sum_{i=1}^N \chi_{[0,\infty)}(u_i) \text{ for } \theta = 0, \tag{3}$$

where $\chi_{[0,\infty)}$ denotes the characteristic functional of $[0, \infty)$ and $\theta \ge 0$ stands for absolute temperature. The remaining quadratic part is characterized by a symmetric, negative semi-definite interaction matrix K. In our computations, we chose $K = \theta_c N(1 - \delta_{ij})_{i,j=1}^N$ (Kronecker–δ) with normalized critical temperature $\theta_c = 1$, i.e., we assume equal interaction of all components and that no self-interaction occurs.

We present fast solvers for discretized Allen-Cahn and Cahn–Hilliard systems that provide global convergence, show robust multigrid convergence speed in numerical computations (with respect to mesh size, temperature $\theta \ge 0$, and the number of phases N), and do not involve any kind of regularization.

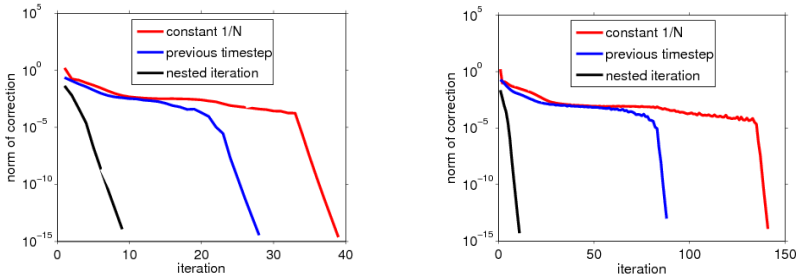

FIGURE 1. Approximate algebraic error over the number of TNNMG iterations for temperature $\theta = 0$ with different initial iterates, mesh size $h_9 = 2^{-8}$ (left) and mesh size $h_{10} = 2^{-9}$ (right).

1. Vector-valued Allen–Cahn system

The L^2-gradient flow of the Ginzburg–Landau free energy $\mathcal{E}$ gives rise to the vector-valued Allen–Cahn system

$$u(t) \in \mathcal{G}: \quad \varepsilon(u_t, v-u) + \varepsilon(\nabla u, \nabla(v-u)) + \tfrac{1}{\varepsilon}(\phi_\theta(v) - \phi_\theta(u)) + \tfrac{1}{\varepsilon}(Ku, v-u) \geq 0 \quad \forall v \in \mathcal{G}$$

with Gibbs constraints $\mathcal{G} = \{v \in H^1(\Omega)^N \mid v(x) \in G \text{ a.e. in } \Omega\}$ and the proper convex, lower semi-continuous functional $\phi_\theta(v) = \int_\Omega \Phi_\theta(v)$. We emphasize that this formulation makes sense for all $\theta \geq 0$. For existence and uniqueness results we refer, e.g., to [3]. After implicit Euler discretization with uniform step size $\tau > 0$ satisfying the stability constraint $\tau < -\varepsilon^2/\lambda_{\min}(K)$ and piecewise linear finite elements $\mathcal{S}_j^N \subset H^1(\Omega)^N$ with respect to a simplicial partition $\mathcal{T}_j$ with set of vertices $\mathcal{N}_j$ and mesh size $h_j = \mathcal{O}(2^{-j})$, a minimization problem of the form

$$u_j \in \mathcal{G}_j: \qquad \mathcal{J}(u_j) \leq \mathcal{J}(v) \qquad \forall v \in \mathcal{G}_j \tag{4}$$

has to be solved in each time step. Here, we impose discrete Gibbs constraints $\mathcal{G}_j = \{v \in \mathcal{S}_j^N \mid v(p) \in G \quad \forall p \in \mathcal{N}_j\}$, $\mathcal{J}(v) = \frac{\tau}{2}a(v,v) + \frac{\tau}{\varepsilon^2}\phi_\theta^{\mathcal{T}}(v) - (u^{\text{old}}, v)$ is a strictly convex energy functional, involving the bounded, coercive bilinear form $a(u,v) = ((I + \frac{\tau}{\varepsilon^2}K)u, v) + \tau(\nabla u, \nabla v)$, a lumped approximation $\phi_\theta^{\mathcal{T}}$ of ϕ_θ, and the solution u^{old} from the previous time step. For the fast and robust algebraic solution of (4) we suggest a novel multigrid method which can be regarded as a variant of polygonal Monotone Multigrid [7] without coarse grid constraints and analyzed as a polygonal version of Truncated Non-smooth Newton Multigrid (TNNMG) [4]. Figure 1 shows the iteration histories for a $V(1,1)$ cycle applied to the problem in the first time step with $N = 4$ components, $\theta = 0$ (deep quench limit), $\varepsilon^2 = 1.6 \cdot 10^{-3}$, $\tau = \frac{3}{4N}\varepsilon^2$, and mesh size $h_j = 2^{-j+1}$ as obtained by $j = 9, 10$ uniform refinements of the initial partition $\mathcal{T}_0$ of the computational domain $\Omega = (-1,1)^2$ consisting of two congruent triangles. Comparing $j = 9$ (left) with $j = 10$ (right) refinement steps, we observe a mesh-dependent transient

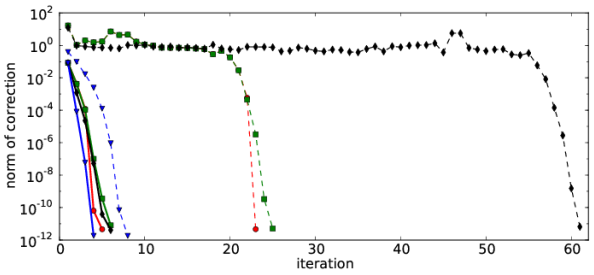

FIGURE 2. Approximate algebraic error over the number of NSNMG iterations for the temperatures $\theta = 0.0$ ($\circ$), 0.001 ($\square$), 0.1 ($\Diamond$), 0.5 ($\triangledown$) with initial iterates zero (dotted lines) and nested iteration (solid lines) for the first spatial problem.

phase for "bad" initial iterates (constant $1/N$ or previous time step). However, textbook multigrid convergence is achieved by nested iteration.

2. Multicomponent Cahn–Hilliard systems

For positive temperature $\theta > 0$, the multicomponent Cahn-Hilliard system associated with the energy $\mathcal{E}$ defined in (1) amounts to find $u, w \in H^1(\Omega)^N$ satisfying

$$(u_t, v) + (L\nabla w, \nabla v) = 0, \quad \varepsilon^2(\nabla u, \nabla v) + (P\Psi_\theta'(u), v) - (w, v) = 0 \quad \forall v \in H^1(\Omega)^N$$

with a positive semi-definite matrix L with kernel spanned by $\mathbf{1}$, the projection $P = I - \frac{1}{N}(\mathbf{1}, \dots, \mathbf{1}) \in \mathbb{R}^{N\times N}$, homogeneous Neumann boundary conditions and suitable initial data. As a starting point for robust numerical solution, we derive a variational inequality formulation that is meaningful for all $\theta \geq 0$ by replacing the chemical potential w by its projection Pw. After semi-implicit discretization in time [1, 2], finite element discretization in space and weak formulation of the algebraic constraints $u \cdot \mathbf{1} = 1$ (see [6]), we arrive at the discrete problem to find $u \in \mathcal{S}^N, w \in \mathcal{S}^N$ such that

$$\begin{aligned} \varepsilon^2(\nabla u, \nabla(v-u)) + \phi_\theta^{\mathcal{T}}(v) - \phi_\theta^{\mathcal{T}}(u) - (w, v-u) &\geq -(Ku^{\mathrm{old}}, v-u) \quad \forall v \in \mathcal{S}^N, \\ -(u, v) - \tau(L\nabla w, \nabla v) &= -(u^{\mathrm{old}}, v) \qquad \forall v \in \mathcal{S}^N. \end{aligned}$$

This formulation allows a direct application of Nonsmooth Schur–Newton Multigrid (NSNMG) methods which were first introduced for saddle point problems with obstacles [5], and later extended to more general free energies including the logarithmic potential [4]. Figure 2 (taken from [6]) shows the iteration history for the first time step, $N = 4$ components, $\varepsilon^2 = 5 \cdot 10^{-3}$, $\tau = 10^{-3}$, and mesh size $h_8 = 2^{-7}$ as obtained by $j = 8$ uniform refinements. While a mesh- and temperature-dependent transient phase is clearly visible for "bad" initial iterate zero, robust textbook multigrid convergence is achieved by nested iteration.

References

[1] J.F. Blowey, M.I.M. Copetti, and C.M. Elliott. Numerical analysis of a model for phase separation of a multi-component alloy. *IMA J. Numer. Anal.*, 16(1):111–139, 1996.
[2] J.F. Blowey and C.M. Elliott. The Cahn-Hilliard gradient theory for phase separation with non-smooth free energy part II: Numerical analysis. *Europ. J. Appl. Math.*, 3:147–179, 1992.
[3] L. Bronsard and F. Reitich. *On three-phase boundary motion and the singular limit of a vector-valued Ginzburg-Landau equation. Arch. Ration. Mech. Anal.*, 124(4):355-379, 1993.
[4] C. Gräser. *Convex Minimization and Phase Field Models.* PhD thesis, FU Berlin, 2011.
[5] C. Gräser and R. Kornhuber. Nonsmooth Newton methods for set-valued saddle point problems. *SIAM J. Numer. Anal.*, 47(2):1251–1273, 2009.
[6] C. Gräser, R. Kornhuber, and U. Sack, *Nonsmooth Schur-Newton methods for multicomponent Cahn-Hilliard systems*, IMA J. Numer. Anal. doi: 10.1093/imanum/dru014 (2014).
[7] R. Kornhuber and R. Krause. *Robust multigrid methods for vector-valued Allen-Cahn equations with logarithmic potential Comp. Vis. Sc.*, 9:103-116, 2006.

Parallel In-Time Methods for Large Scale Problems

Rolf Krause
(joint work with D. Ruprecht, R. Speck, M. Bolten)

Computational science, like other disciplines using high-performance computing, faces enormous challenges from the exponential increase in the number of cores in state-of-the-art HPC systems. Utilizing such machines efficiently not only requires sophisticated implementation techniques but also the use of mathematical algorithms inherently suited for parallelization: Even a perfectly optimized implementation of an algorithm with strong serial dependencies will not be able to fully exploit the million-way concurrency at hand. As a consequence, the computer architectures available nowadays require the development of discretization and solution algorithms which are capable of exploiting efficiently massive parallelism.

Traditionally, parallelization for the solution of time dependent partial differential equations is usually done "in space" In fact, for the numerical solution of time-dependent partial differential equations, many approaches and studies exist that investigate strategies for parallelization in space. In the temporal direction however, most available methods (e.g. Runge-Kutta, multi-step methods, etc.) employ a form of time-stepping, where step after step is computed in a strictly sequential fashion. In some sense, these classical time-marching schemes constitute a bottleneck: If spatial and temporal resolution of a discretization both are increased, the computational cost **per time step** can be kept constant by employing more processors for the spatial direction (assuming good weak scaling, that is). The increased cost of having to perform **more** time steps however cannot be mitigated. By offering parallelism in the temporal direction in addition to spatial parallelization, time-parallel methods can increase concurrency in numerical codes and help to address this issue and to better utilize massively parallel systems.

It therefore seems natural to develop solution algorithms, which are also parallel in time. At a first glance, this seems to be counterintuitive, as parallelization

in time seems to be in strong contrast with the evolutionary character of an initial value problem. However, as it turned out, using well balanced iterative or predictor-corrector strategies, also parallelization in time is possible - although currently still at the price of reduced parallel efficiency.

In this talk, different iterative methods for the solution of time dependent PDEs, which provide concurrency along the temporal axis, are presented, discussed, and evaluated. We will start from the by now well established Parareal, see [1] for a survey or [2], Parareal relies on replacing step-by-step integration in time with an iterative, multi-level procedure: solutions on different levels of accuracy are computed and combined, where the main computational cost on the fine levels is parallelized over multiple time-slices. Usually, the efficiency depends critically on the number of iterations required: If convergence is too slow, time-parallel methods can actually become slower than serial time-stepping.

We therefor discuss the more complex "parallel full appoximation scheme in space and time" (PFASST), see [3]. PFASST is build on using Spectral Deferred Correction (SDC) methods on multiple levels *in time*. SDC methods are build on the reformulation of the initial value problem on each of the considered time slices in terms of a Picard integral equation. The advantage of this formulation is that it allows for approximating the sought trajectory in an iterative manner (SDC-sweeps). By interpreting this procedure as a preconditioned iterative solution method, we can investigate its convergence properties for Dahlquist's test equation. We then extend this convergence analysis to the case of multiple levels in time (MLSDC), where the SDC iteration is accelerated by additional SDC-sweeps on coarse time levels. As this analysis shows, the coarse SDC sweeps do not accelerate the convergence significantly, but at least allow for the creation of a fast propagator for the initial value on a much coarser mesh in time [5].

We then proceed to PFASST, which exploits the idea of multilevel SDC in order to create a time-parallel method which is of much higher efficiency as Parareal. The main advantage of PFASST is that it distributes the work for approximating the trajectory on the finest level in time over several levels, which allows to increase the parallelism.

We discuss several factors that govern the performance of the methods and show examples of good and bad convergence.

We discuss the properties of the resulting space/time parallel method and present recent performance results on large-scale parallel machines shown that demonstrate how parallelization in time can achieve good efficiency if a suitably optimized coarse level representation is used. One example illustrates how PFASST can successfully improve strong scaling in a run using all 448K cores of the IBM BlueGene/Q installation JUQUEEN, see [6] and also [4].

References

[1] M. J. Gander and S. Vandewalle, *Analysis of the parareal time-parallel time-integration method*, SIAM Journal on Scientific Computing **29** (2007), 556–578.

[2] M. Gander and E. Hairer, *Nonlinear convergence analysis for the parareal algorithm*, Domain Decomposition Methods in Science and Engineering **60** (2008), pp. 45–56.

[3] M. Emmett and M. L. Minion, *Toward an efficient parallel in time method for partial differential equations*, Communications in Applied Mathematics and Computational Science **7** (2012), pp. 105–132.

[4] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. Minion, M. Winkel, and P. Gibbon, *A massively space-time parallel N-body solver*, Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, **SC '12** (2012), pp. 92:1–92:11.

[5] R. Speck, D. Ruprecht, M. Emmett, M. Minion, M. Bolten, and R. Krause, *A multi-level spectral deferred correction method*, arXiv:1307.1312 (2014). (submitted to BITN, in revision)

[6] D. Ruprecht, R. Speck, M. Emmett, M. Bolten, and R. Krause, *Poster: Extreme-scale space-time parallelism*, Proceedings of the 2013 Conference on High Performance Computing Networking, Storage and Analysis Companion **SC '13 Companion** (2013).

Applications for Filtering Algebraic Multigrid

Arne Nägel
(joint work with Gabriel Wittum)

With density driven flow and poroelasticity we consider two problem classes of coupled PDEs arising in geosciences. Both problems are differential algebraic equations. We comment on similarities and differences, and outline how the equations can be solved using decoupling and fully-coupled methods, and fine-tuned filtering algebraic multigrid solvers.

The first part of the talk focuses on decoupling strategies, in particular for the nonlinear density driven flow problems. To that end different non-linear solvers have been investigated: In a test case, an iterative coupling (single step nonlinear Gauss-Seidel) outperforms a partial Newton method (single step non-linear Jacobi). In the latter case corrections are computed independently, thus, this scheme shows to be inappropriate for this highly nonlinear problem class, and suffers from severe time-step restrictions. This effect is less pronounced for the iterative coupling, and mitigates after the first simulation phase when the velocity profile has stabilized. The fully coupled Newton method, however, proofs being superior to both iterative approaches, but also requires a monolithic linear solver.

In the second part of the talk, we deal with poroelasticity problems. Owing to the linear character, these systems can be decoupled efficiently by iterative coupling. The mechanics sub-system is time-independent, wthus the problem perfectly suited for algebraic multigrid solvers. A preliminary scalability study (up to 256 cores) shows good scalability.

Finite element error estimation for a Schrödinger operator with inverse-square potential

Jeffrey S. Ovall
(joint work with H. Li)

We develop an a posteriori error estimate of hierarchical type for source and eigenvalue problems for operators of the form $-\Delta + c^2/r^2$ in $\Omega \subset \mathbf{R}^2$, where $c \geq 0$, and $r = |x|$ is the distance to the origin, which is assumed to be in $\overline{\Omega}$. The solutions of such problems generically have an r^c-type singularity at the origin, in addition to the usual geometric singularities at boundary corners. This work is a combination of the results from the papers [1, 2]. For a source problem with solution u, the discrete solution u_n is computed on a conforming mesh $\mathcal{T}_n$ in the space of affine elements which boundary and at the origin, and an approximate error function $\varepsilon_n \approx u - u_n$ is computed by projecting the error onto the space of quadratic elements which vanish on the boundary and at every vertex in the mesh. It is shown that

$$\frac{\|\varepsilon_n\|_0}{\|u - \Pi_n u\|_0} \to 1 \quad , \quad \frac{\|\varepsilon_n\|_1}{\|u - u_n\|_1} \to 1 \quad , \quad \frac{|\!|\!|\varepsilon_n|\!|\!|}{|\!|\!|u - u_n|\!|\!|} \to 1 \quad , \quad \frac{\sum_{T\in\mathcal{T}_n} |\varepsilon_n|_{2,1,T}}{|u|_{2,1}} \to 1$$

on a sequence of geometrically graded meshes $\mathcal{T}_n$. In words, the approximate error function ε_n provides: asymptotically exact estimation of (nodal) interpolation error in L^2, asymptotically exact estimation of discretization error in H^1 and energy norm, and convergent approximation of the Hessian of u in $W^{2,1}$. Although these results are proven on specially graded meshes, numerical experiments indicate that they apply on meshes which are adaptively refined using local indicators $|\!|\!|\varepsilon_n|\!|\!|_T$ to mark triangles. The results for source problems are readily adapted for eigenvalue problems, with analogous results asymptotic exactness results for eigenvalues and eigenfunctions.

References

[1] H. Li and J.S. Ovall, *A posteriori error estimation of hierarchical type for the Schrödinger operator with inverse-square potential*, Numerische Mathematik (2014), DOI : 10.1007/s00211-014-0628-y.

[2] H. Li and J.S. Ovall, *A posteriori eigenvalue error estimation for the Schrödinger operator with inverse-square potential*, Discrete and Continuous Dynamical Systems–Series B (2014), to appear.

On a highly scalable infrastructure for massively parallel multigrid solvers

Sebastian Reiter
(joint work with Andreas Vogel, Gabriel Wittum)

Application of parallel geometric multigrid solvers on adaptively refined grids requires a careful design of the involved load-balancing and load-migration routines as well as fast communication between copies of distributed objects. In error-estimation based refinement strategies, multiple rebalancing steps may be required

to provide a uniform element distribution between all processors in all steps. In order to perform such operations efficiently on supercomputers with millions of cores, one has to extend existing load balancing and communication schemes.

We outline the parallel infrastructure for distributed multigrid hierarchies in the simulation framework UG4 and present experimental scaling studies for our geometric multigrid solver on adaptive and non-adaptive grid hierarchies on up to 262144 processes. We focus on an efficient hierarchical organization of the involved processes, on efficient horizontal and vertical communication schemes as well as on the parallel refinement and load-balancing strategies used.

Solving PDEs on evolving surfaces: An Eulerian space-time FEM

Arnold Reusken
(joint work with Maxim Olshanskii, Jörg Grande, Xianmin Xu)

1. Introduction

Partial differential equations posed on evolving surfaces appear in a number of applications. Recently, several numerical approaches for handling such type of problems have been introduced. We refer to [1] for a recent overview of numerical methods. In this report we outline an Eulerian finite element method studied in [2, 4, 3]. The key idea of this method is to use restrictions of (usual) space-time volumetric finite element functions to the space-time manifold. This trace finite element technique was introduced for stationary surfaces in [5].

In this report we address the key ideas of this space-time trace-FEM and some main results of the error analysis, in particular a result on second order accuracy of the method in space and time. For details we refer to [2, 3]. Results of numerical experiments show that the method is extremely robust and that even for the case with a topological singularity (droplet collision) accurate results can be obtained on a fixed Eulerian (space-time) grid with a large time step.

As a model problem we use the following one. Consider a surface $\Gamma(t)$ passively advected by a *given* smooth velocity field $\mathbf{w} = \mathbf{w}(x,t)$, i.e. the normal velocity of $\Gamma(t)$ is given by $\mathbf{w} \cdot \mathbf{n}$, with $\mathbf{n}$ the unit normal on $\Gamma(t)$. We assume that for all $t \in [0,T]$, $\Gamma(t)$ is a hypersurface that is closed ($\partial\Gamma = \emptyset$), connected, oriented, and contained in a fixed domain $\Omega \subset \mathbb{R}^d$, $d = 2, 3$. In the remainder we consider $d = 3$, but all results have analogs for the case $d = 2$. The convection-diffusion equation on the surface that we consider is given by:

$$\dot{u} + (\mathrm{div}_\Gamma \mathbf{w})u - \nu_d \Delta_\Gamma u = f \qquad \text{on } \Gamma(t),\ t \in (0,T], \tag{1}$$

with a prescribed source term $f = f(x,t)$ and homogeneous initial condition $u(x,0) = u_0(x) = 0$ for $x \in \Gamma_0 := \Gamma(0)$. Here $\dot{u} = \frac{\partial u}{\partial t} + \mathbf{w} \cdot \nabla u$ denotes the advective material derivative, $\mathrm{div}_\Gamma := \mathrm{tr}\left((I - \mathbf{n}\mathbf{n}^T)\nabla\right)$ is the surface divergence and Δ_Γ is the Laplace-Beltrami operator, $\nu_d > 0$ is the constant diffusion coefficient. If we take $f = 0$ and an initial condition $u_0 \neq 0$, this surface PDE is obtained from mass conservation of the scalar quantity u with a diffusive flux

on $\Gamma(t)$. A standard transformation to a homogeneous initial condition, which is convenient for a theoretical analysis, leads to (1).

2. Well-posed space-time weak formulation

Several weak formulations of (1) are known in the literature. The most appropriate for our purposes is a integral space-time formulation proposed in [2]. In this section we outline this formulation. Consider the space-time manifold

$$\mathcal{S} = \bigcup_{t\in(0,T)} \Gamma(t)\times\{t\}, \quad \mathcal{S}\subset\mathbb{R}^4.$$

On $L^2(\mathcal{S})$ we use the scalar product $(v,w)_0 = \int_0^T \int_{\Gamma(t)} vw\, ds\, dt$. Let ∇_Γ denote the tangential gradient for $\Gamma(t)$ and introduce the space

$$H = \{\, v\in L^2(\mathcal{S}) \mid \|\nabla_\Gamma v\|_{L^2(\mathcal{S})} < \infty \,\},\ (u,v)_H := (u,v)_0 + (\nabla_\Gamma u, \nabla_\Gamma v)_0. \tag{2}$$

We consider the material derivative $\dot{u}$ of $u\in H$ as a distribution on $\mathcal{S}$:

$$\langle \dot{u},\phi\rangle = -\int_0^T \int_{\Gamma(t)} u\dot{\phi} + u\phi\, \mathrm{div}_\Gamma \mathbf{w}\, ds\, dt \quad \text{for all } \phi\in C_0^1(\mathcal{S}).$$

In [2] it is shown that $C_0^1(\mathcal{S})$ is dense in H. If $\dot{u}$ can be extended to a bounded linear functional on H, we write $\dot{u}\in H'$. Define the space

$$W = \{\, u\in H \mid \dot{u}\in H' \,\}, \quad \text{with } \|u\|_W^2 := \|u\|_H^2 + \|\dot{u}\|_{H'}^2.$$

In [2] properties of H and W are derived. Both spaces are Hilbert spaces and smooth functions are dense in H and W. Define

$$\overset{\circ}{W} := \{\, v\in W \mid v(\cdot,0) = 0 \quad \text{on } \Gamma_0 \,\}.$$

The space $\overset{\circ}{W}$ is well-defined, since functions from W have well-defined traces in $L^2(\Gamma(t))$ for any $t\in[0,T]$. We introduce the symmetric bilinear form

$$a(u,v) = \nu_d(\nabla_\Gamma u, \nabla_\Gamma v)_0 + (\,\mathrm{div}_\Gamma \mathbf{w}\, u, v)_0, \quad u,v\in H,$$

which is continuous on $H\times H$. The weak space-time formulation of (1) reads: For given $f\in L^2(\mathcal{S})$ find $u\in \overset{\circ}{W}$ such that

$$\langle \dot{u}, v\rangle + a(u,v) = (f,v)_0 \quad \text{for all } v\in H. \tag{3}$$

In [2] the inf-sup property

$$\inf_{0\neq u\in \overset{\circ}{W}} \sup_{0\neq v\in H} \frac{\langle \dot{u},v\rangle + a(u,v)}{\|u\|_W \|v\|_H} \geq c_s > 0 \tag{4}$$

is proved. Using this in combination with the continuity result one can show that the *weak formulation* (3) *is well-posed.*

We introduce a similar "*time-discontinuous*" weak formulation that is better suited for the finite element method that we consider. We take a partitioning of the time interval: $0 = t_0 < t_1 < \ldots < t_N = T$, with a (for simplicity0 uniform time step $\Delta t = T/N$. A time interval is denoted by $I_n := (t_{n-1}, t_n]$. The symbol $\mathcal{S}^n$

denotes the space-time interface corresponding to I_n, i.e., $\mathcal{S}^n := \cup_{t\in I_n}\Gamma(t)\times\{t\}$, and $\mathcal{S} := \cup_{1\le n\le N}\mathcal{S}^n$. We introduce the following subspaces of H:

$$H_n := \{\, v \in H \mid v = 0 \quad \text{on } \mathcal{S}\setminus\mathcal{S}^n \,\},$$

and define the spaces

$$W_n = \{\, v \in H_n \mid \dot v \in H_n' \,\}, \quad \|v\|^2_{W_n} = \|v\|^2_H + \|\dot v\|^2_{H_n'}, \tag{5}$$

$$W^b := \oplus_{n=1}^N W_n, \text{ with norm } \|v\|^2_{W^b} = \sum_{n=1}^N \|v\|^2_{W_n}. \tag{6}$$

For $u \in W_n$, the one-sided limits $u_+^n = u_+(\cdot,t_n)$ (i.e., $t \downarrow t_n$) and $u_-^n = u_-(\cdot,t_n)$ (i.e., $t \uparrow t_n$) are well-defined in $L^2(\Gamma(t_n))$. At t_0 and t_N only u_+^0 and u_-^N are defined. For $v \in W^b$, a jump operator is defined by $[v]^n = v_+^n - v_-^n \in L^2(\Gamma(t_n))$, $n = 1,\dots,N-1$. For $n = 0$, we define $[v]^0 = v_+^0$. On the cross sections $\Gamma(t_n)$, $0 \le n \le N$, of $\mathcal{S}$ the L^2 scalar product is denoted by $(\psi,\phi)_{t_n} := \int_{\Gamma(t_n)} \psi\phi\, ds$. In addition to $a(\cdot,\cdot)$, we define on the broken space W^b the bilinear forms

$$d(u,v) = \sum_{n=1}^N d^n(u,v), \quad d^n(u,v) = ([u]^{n-1}, v_+^{n-1})_{t_{n-1}}, \quad \langle \dot u, v\rangle_b = \sum_{n=1}^N \langle \dot u_n, v_n\rangle\,.$$

One can show that the unique solution to (3) is also the unique solution of the following variational problem in the broken space: Find $u \in W^b$ such that

$$\langle \dot u, v\rangle_b + a(u,v) + d(u,v) = (f,v)_0 \quad \text{for all } v \in W^b. \tag{7}$$

For this time discontinuous weak formulation an inf-sup stability result (that is weaker than the one in (4)) can be derived. From an algorithmic point of view the formulation (7) has the advantage that due to the use of the broken space $W^b = \oplus_{n=1}^N W_n$ it can be solved in a time stepping manner.

3. Space-time finite element method

We introduce a finite element method which is a Galerkin method with $W_h \subset W^b$ applied to the variational formulation (7). To define this W_h, consider the partitioning of the space-time volume domain $Q = \Omega\times(0,T] \subset \mathbb{R}^{3+1}$ into time slabs $Q_n := \Omega\times I_n$. Corresponding to each time interval $I_n := (t_{n-1},t_n]$ we assume a given shape regular tetrahedral triangulation $\mathcal{T}_n$ of the spatial domain Ω. The corresponding spatial mesh size parameter is denoted by h. Then $\mathcal{Q}_h = \bigcup_{n=1,\dots,N} \mathcal{T}_n\times I_n$ is a subdivision of Q into space-time prismatic nonintersecting elements. We shall call $\mathcal{Q}_h$ a space-time triangulation of Q. Note that this triangulation is not necessarily fitted to the surface $\mathcal{S}$. We allow $\mathcal{T}_n$ to vary with n.

For any $n \in \{1,\dots,N\}$, let V_n be the finite element space of continuous piecewise linear functions on $\mathcal{T}_n$. We define the *volume space-time finite element space*:

$$V_h := \{\, v : Q \to \mathbb{R} \mid v(x,t) = \phi_0(x) + t\phi_1(x) \text{ on every } Q_n, \text{ with } \phi_0,\, \phi_1 \in V_n \,\}. \tag{8}$$

Thus, V_h is a space of piecewise P1 functions with respect to $\mathcal{Q}_h$, continuous in space and discontinuous in time. Now we define our *surface finite element space* as the space of traces of functions from V_h on $\mathcal{S}$:

$$W_h := \{\, w : \mathcal{S} \to \mathbb{R} \mid w = v_{|\mathcal{S}},\ v \in V_h \,\}. \tag{9}$$

The finite element method reads: Find $u_h \in W_h$ such that

$$\langle \dot{u}_h, v_h \rangle_b + a(u_h, v_h) + d(u_h, v_h) = (f, v_h)_0 \quad \text{for all } v_h \in W_h. \tag{10}$$

Due to $u_h \in H^1(Q_n)$ for all $n = 1, \ldots, N$, the first term in (10) can be written as

$$\langle \dot{u}_h, v_h \rangle_b = \sum_{n=1}^{N} \int_{t_{n-1}}^{t_n} \int_{\Gamma(t)} (\frac{\partial u_h}{\partial t} + \mathbf{w} \cdot \nabla u_h) v_h ds\, dt.$$

The method can be implemented with a time marching strategy. Of course, for the implementation of the method one needs a quadrature rule to approximate the integrals over $\mathcal{S}^n$. Implementation aspects of the method are discussed in [4].

4. Discretization error analysis

In this section we briefly address the discretization error analysis of the method (10), which is presented in [3]. We first explain a discrete mass conservation property of the scheme (10). We consider the case that (1) is derived from mass conservation of a scalar quantity with a diffusive flux on $\Gamma(t)$. The original problem then has a nonzero initial condition u_0 and a source term $f \equiv 0$. The solution u of the original problem has the mass conservation property $\bar{u}(t) := \int_{\Gamma(t)} u\, ds = \int_{\Gamma(0)} u_0\, ds$ for all $t \in [0, T]$. After a suitable transformation one obtains the equation (1) with a zero initial condition u_0 and a right hand-side f which satisfies $\int_{\Gamma(t)} f\, ds = 0$ for all $t \in [0, T]$. The solution u of (1) then has the "shifted" mass conservation property $\bar{u}(t) = 0$ for all $t \in [0, T]$. The discrete problem (10) has the following weaker mass conservation property, with $\bar{u}_h(t) := \int_{\Gamma(t)} u_h\, ds$:

$$\bar{u}_{h,-}(t_n) = 0 \quad \text{and} \quad \int_{t_{n-1}}^{t_n} \bar{u}_h(t)\, dt = 0, \quad n = 1, 2, \ldots N. \tag{11}$$

For a *stationary* surface, $\bar{u}_h(t)$ is a piecewise affine function and thus (11) implies $\bar{u}_h(t) \equiv 0$, i.e,. we have exact mass conservation on the discrete level. If the surface evolves, the finite element method is not necessarily mass conserving: (11) holds, but $\bar{u}_h(t) \neq 0$ may occur for $t_{n-1} \leq t < t_n$. In the discretization error analysis we use a *consistent* stabilizing term involving the quantity $\bar{u}_h(t)$:

$$a_\sigma(u, v) := a(u, v) + \sigma \int_0^T \bar{u}(t)\bar{v}(t)\, dt, \quad \sigma \geq 0. \tag{12}$$

Instead of (10) we consider the stabilized version: Find $u_h \in W_h$ such that

$$\langle \dot{u}_h, v_h \rangle_b + a_\sigma(u_h, v_h) + d(u_h, v_h) = (f, v_h)_0 \quad \text{for all } v_h \in W_h. \tag{13}$$

Taking $\sigma > 0$ we expect both a stabilizing effect and an improved discrete mass conservation property. Ellipticity of the bilinear form and error bounds are derived in the mesh-dependent norm:

$$|\!|\!| u |\!|\!|_h := \left(\|u_-^N\|_T^2 + \sum_{n=1}^{N} \|[u]^{n-1}\|_{t_{n-1}}^2 + \|u\|_H^2 \right)^{\frac{1}{2}}.$$

In the error analysis we need the following condition: there exists a $c_0 > 0$ such that

$$\operatorname{div}_\Gamma \mathbf{w}(x,t) + \nu_d c_F(t) \geq c_0 \quad \text{for all } x \in \Gamma(t),\ t \in [0,T]. \tag{14}$$

Here $c_F(t) > 0$ results from the Poincare inequality

$$\int_{\Gamma(t)} |\nabla_\Gamma u|^2 \, ds \geq c_F(t) \int_{\Gamma(t)} (u - \frac{1}{|\Gamma(t)|}\bar{u})^2 \, ds \quad \forall\ t \in [0,T],\ \forall\ u \in H. \tag{15}$$

A main result derived in [3] is given in the following theorem. We assume that the time step Δt and the spatial mesh size parameter h have comparable size: $\Delta t \sim h$.

Theorem 1. *Assume* (14) *and take* $\sigma \geq \frac{\nu_d}{2} \max\limits_{t\in[0,T]} \frac{c_F(t)}{|\Gamma(t)|}$, *where* $c_F(t)$ *is defined in* (15). *Then the ellipticity estimate*

$$\langle \dot{u}, u \rangle_b + a_\sigma(u,u) + d(u,u) \geq c_s |\!|\!| u |\!|\!|_h^2 \quad \textit{for all } u \in W^b$$

holds, with $c_s = \frac{1}{2}\min\{1, \nu_d, c_0\}$ *and* c_0 *from* (14). *Let* $u \in \overset{\circ}{W}$ *be the solution of* (3) *and assume* $u \in H^2(\mathcal{S})$. *For the solution* $u_h \in W_h$ *of the discrete problem* (13) *the following error bound holds:*

$$|\!|\!| u - u_h |\!|\!|_h \leq ch \|u\|_{H^2(\mathcal{S})}.$$

A further main result derived in [3] is related to second order convergence. Denote by $\|\cdot\|_{-1}$ the norm dual to the $H_0^1(\mathcal{S})$ norm with respect to the L^2-duality. Under the conditions given in Theorem 1 and some further mild assumptions the error bound

$$\|u - u_h\|_{-1} \leq ch^2 \|u\|_{H^2(\mathcal{S})}$$

holds.

For results of numerical experiments with this space-time method we refer to [2, 4].

References

[1] G. Dziuk, C.M Elliott, *Finite element methods for surface PDEs*, Acta Numerica **22** (2013), 289–396.

[2] M. Olshanskii, A. Reusken, X. Xu, *An Eulerian space-time finite element method for diffusion problems on evolving surfaces* Preprint University of Houston, NA&SC Preprint **5** (2013). Accepted for publication in SIAM J. Numer. Anal. (2014).

[3] M. Olshanskii, A. Reusken, *Error analysis of a space-time finite element method for solving PDEs on evolving surfaces*, Preprint RWTH Aachen University, IGPM Preprint **376** (2013). Accepted for publication in SIAM J. Numer. Anal. (2014).

[4] J. Grande, *Finite element methods for parabolic equations on moving surfaces*, Preprint RWTH Aachen University, IGPM Preprint **360** (2013). Accepted for publication in SIAM J. Sci. Comput. (2014).
[5] M. Olshanskii, A. Reusken, J. Grande, *A finite element method for elliptic equations on surfaces*, SIAM J. Numer. Anal. **47** (2009), 3339–3358.

A posteriori error estimator for nonlinear interface conditions discretized by Finite Volumes – Applications in Neuronscience

Gillian Queisser
(joint work with Markus Breit)

In this talk we introduce a broad range of neuroscientific questions that are addressed using numerical methods. Based on spatially resolved complex neuronal morphologies, models and numerical methods are introduced in order to investigate structure-function interplay at the cellular level. This leads to a general model, consisting of coupled systems of non-linear reaction-diffusion type equations, with highly non-linear interface conditions. The analysis of these equations show that adaptive methods, based on Finite Volumes, can greatly reduce computational cost. Here we introduce a novel a posteriori error estimator for the coupled pde-system under consideration and give an outline of the proof. We summarize by applying the a posteriori error estimator for adaptive grid refinement to different biologically motivated case studies.

On the Foundation of Efficient Preconditioners for PDE Constrained Shape Optimization Problems

Volker Schulz
(joint work with Martin Siebenborn, Kathrin Welker)

Shape optimization is a very active field of research with numerous applications of economic importance. Several examples from aerodynamics, acoustics and thermoelastics are used to illustrate this and to motivate the following more theoretical considerations. Although parametric geometry description (like CAD) are widely used in industry, they lead to high numerical costs for non-trivial geometry resolutions and pose severe limitations to the set of reachable shapes. The alternative avoiding these problems is the nonparametric approach which leaves all mesh nodes describing the geometry under investigation free for optimization and is based on the shape calculus as surveyed, e.g., in the monographies [1, 2]. The current numerical state of the art in shape optimization based on the shape calculus is characterized by first order methods of steepest descent type and a general lack of second order methods. However, ideas from second order methods aiming a Newton-like strategies give rise to excellent preconditioners as demonstrated in [3, 4].

In this talk, a general framework based on the differential geometric investigations in [5] is presented, which considers the set of admissible shapes as a Riemannian manifold and constructs Taylor series expansion and Newton methods similar to the ideas in [6] for finite dimensional matrix manifold. This novel approach, published in [7], introduces a Riemannian shape Hessian as a Hessian formulation for second shape derivatives which, in contrast to the second shape derivative (which is so far historically but misleadingly named shape Hessian) possesses the properties which are expected from a Hessian: symmetry and provision of a Taylor series expansion. This approach is carried on in [8] to PDE constrained shape optimization and develops a novel sequential quadratic programming framework for shape optimization based on the shape calculus, where the linear-quadratic subproblems to be solved in each nonlinear iteration have the structure of usual optimal control problems and are thus accessible to the wealth of efficient methods developed for this problem class like, e.g., multigrid optimization methods [9, 10].

References

[1] J. Sokolowski and J.-P. Zolésio. *Introduction to Shape Optimization: Shape Sensitivity Analysis*. Springer, 1992.

[2] M. C. Delfour and J.-P. Zolésio. *Shapes and Geometries: Analysis, Differential Calculus, and Optimization*. Advances in Design and Control. SIAM Philadelphia, 2001.

[3] S. Schmidt. *Efficient Large Scale Aerodynamic Design Based on Shape Calculus*. PhD thesis, Universität Trier, 2010.

[4] S. Schmidt, C. Ilic, V. Schulz, and N. Gauger. Three-dimensional large-scale aerodynamic shape optimization based on shape calculus. *AIAA Journal*, 51(11):2615–2627, 2013.

[5] P. Michor and D. Mumford. Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc. (JEMS)*, 8:1–48, 2006.

[6] P. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.

[7] V. Schulz. A Riemannian view on shape optimization. *Foundations of Computational Mathematics*, 14:483–501, 2014.

[8] V. Schulz, M. Siebenborn, and K. Welker. Towards a lagrange-newton approach for pde constrained shape optimization. submitted to G. Leugering et al., editors, Trends in PDE Constrained Optimization, volume 165 of International Series of Numerical Mathematics, Springer 2014. (arXiv:1405.3266), 2014.

[9] A. Borzi and V. Schulz. Multigrid methods for PDE optimization. *SIAM Review*, 51(2):361–395, 2009.

[10] A. Borzì and V. Schulz. *Computational optimization of systems governed by partial differential equations*. Number 08 in SIAM book series on Computational Science and Engineering. SIAM Philadelphia, 2012.

A robust Petrov-Galerkin discretisation of convection-diffusion equations

Rob Stevenson
(joint work with Dirk Broersen)

It is well-known that standard Galerkin discretisations of convection-diffusion equations fail to deliver good approximations for a vanishing diffusion term. Therefore, we consider *Petrov-Galerkin* discretisations.

Unless the layers are resolved by the mesh, the H^1-errors of finite element approximations will be dominated by the errors in the layers. This holds also true for L_2-errors when conforming finite elements are applied due to the strong enforcement of Dirichlet boundary conditions. Therefore, we prefer to measure the errors in the L_2-norm, and to allow for discontinuous approximations. To this end, we consider an *ultra-weak* variational formulation of the convection-diffusion equation in mixed form. It is shown to define a boundedly invertible mapping $U \to V'$, with U and V being Hilbert spaces, where U is (essentially) a multiple copy of the L_2-space.

Building on the earlier works [2, 4, 3], we equip V with the operator-dependent *optimal test norm*. Then given a finite dimensional trial space $U^h \subset U$, the Petrov-Galerkin discretisation with the *optimal test space* delivers the best approximation from U^h to the solution w.r.t. the norm on U.

To arrive at an implementable method, this truly optimal test space has to be replaced by its projection onto a finite dimensional *test search space* V^h. With common variational formulations, the truly optimal test functions exhibit layers, and for vanishing diffusion, the test search space has to be chosen increasingly large to get satisfactory results.

We construct a non-standard variational formulation such that for a zero diffusion term, the discrete system is a well-posed Petrov-Galerkin discretisation of the limiting transport problem. This can be seen as a necessary condition for the equations, which define the optimal test functions, not to be singularly perturbed.

Numerical experiments, reported in [1], show that with a fixed test search space V^h, only dependent on U^h and with $\dim V^h \leq C \dim U^h$, the obtained approximations are very close to the *best* approximations to the solution from the trial space, uniformly in the size of the diffusion term.

References

[1] D. Broersen and R. P. Stevenson. A robust Petrov-Galerkin discretization of convection-diffusion equations. *Comput. Math. Appl.*, 2014. To appear.
[2] J. W. Barrett and K. W. Morton. Approximate symmetrization and Petrov-Galerkin methods for diffusion-convection problems. *Comput. Methods Appl. Mech. Engrg.*, 45(1-3), 1984.
[3] A. Cohen, W. Dahmen, and G. Welper. Adaptivity and variational stabilization for convection-diffusion equations. *ESAIM: Mathematical Modelling and Numerical Analysis*, 46:1247–1273, 2012.
[4] L. Demkowicz and J. Gopalakrishnan. A class of discontinuous Petrov-Galerkin methods. II. Optimal test functions. *Numer. Methods Partial Differential Equations*, 27(1):70–105, 2011.

Multigrid for far and near field maps of the Helmholtz equation

Wim Vanroose

Almost all knowledge about microscopically small physical systems such as molecules is obtained through scattering experiments that are described mathematically by a high-dimensional Helmholtz equation. In contrast to engineering problem, where there are material jumps, these systems are characterized by a smoothly varying wave number.

In this talk we show that the far and near field scattering amplitude of a Helmholtz equation with smoothly varying wave numbers can be efficiently be calculated. Indeed, these amplitudes are integral expressions over the solution of the Helmholtz equation solved on a finite numerical box with absorbing boundary conditions covering the object of interest.

A typical calculation then consists out of two steps: First, solve the high-dimensional Helmholtz equation, which is a computationally very expensive and requires supercomputer infrastructure for the most challenging problems. Secondly, integrate over the solution to obtain the far field or near field amplitude. The latter is cheap can be done as post processing step on a laptop.

By deforming the contour of integration into complex plane in step 2, the Helmholtz problem, step 1, becomes much easier. Indeed, we show that the deformation of the contour makes the Helmholtz into a Complex shifted Laplacian problem, which is known to be solvable in a scalable way.

We validate the method for some benchmark problems and show the $\mathcal{O}(n)$ scalability on 3D Helmholtz and Schrödinger equations.

Computing Edge Weights for Community Detection Through Bipartite Matching

Panayot S. Vassilevski

(joint work with Van Emden Henson, David Hysom, Geoffrey Sanders, Andy Yoo)

For a given graph G with a vertex set $V = \{i\}$ and edge set $E = \{e\}$ where each $e = (i, j)$ is a undirected pair of vertices, we define a nonlinear functional

$$J(\mathbf{w}) = \sum_{i \in V} \left(\frac{1}{w_i - w_{\min}} + \frac{1}{w_{\max} - w_i} + \sum_{e=(i,j) \in E} \frac{1}{|w_i - w_j|} \right).$$

We are looking to minimize J for $\mathbf{w} = (w_i)$ where each w_i varies in a given interval $(w_{\min}, w_{\max})$, for example $(0, 1)$. To compute a minimum of J, we use a nonlinear Gauss–Seidel iteration by minimizing for every given vertex i with respect to w_i the local functional

$$J_{loc}(w_i) = \frac{1}{w_i - w_{\min}} + \frac{1}{w_{\max} - w_i} + \sum_{e=(i,j) \in E} \frac{1}{|w_i - w_j|},$$

assuming that the neighboring values w_j, for $e = (i, j)$ are given.

Once a minimum is computed (within certain accuracy), we assign to each edge $e = (i, j)$ the weight $\widehat{\lambda}_e = \frac{1}{|w_i - w_j|}$. A large weight indicates that the vertices i and j are strongly connected. The notion of strength of connectivity quantified by our edge weights is important to design algorithms for community detection (cf. [5]). We go one step further by introducing a rigorous definition of "community".

Based on the computed edge weights, we devise a recursive multilevel algorithm (combining ideas from [1], [3], [4], and [2]) that matches vertices for which a "surface-to-volume" ratio is large. The latter is defined for any pair of aggregates (initially aggregates are single vertices) $\mathcal{A}$ and $\mathcal{A}'$, as follows

$$\eta_{\mathcal{A},\mathcal{A}'} = \frac{\widehat{\lambda}_{\mathcal{A},\,\mathcal{A}'}}{\sqrt{d_{\mathcal{A}} d_{\mathcal{A}'}}},$$

where $d_{\mathcal{A}} = \sum_{e=(i,j):\, i,j\in\mathcal{A}} \widehat{\lambda}_e$ and $\widehat{\lambda}_{\mathcal{A},\,\mathcal{A}'} = \sum_{e=(i,j):\, i\in\mathcal{A},\, j\in\mathcal{A}'} \widehat{\lambda}_e$. For any partition $\pi = \{\mathcal{A}\}$ of nonoverlapping sets of vertices (referred to as aggregates), we define an "energy" functional

$$E_\pi = \sum_{\mathcal{A},\mathcal{A}'} \eta_{\mathcal{A},\mathcal{A}'} |d_{\mathcal{A}} - d_{\mathcal{A}'}|.$$

Finally, we define "community" as a partition π with minimal energy among some other possible partitions.

The above defined edge weights and respective recursive multilevel matching algorithm is successfully tested on some real-life networks with known "ground truth" communities available from the Stanford Large Network Dataset Collection (SNAP) as described in [6].

References

[1] M. Luby, *A simple parallel algorithm for the maximal independent set algorithm,* SIAM Journal on Computing **15**(4)(1986), pp. 1036-1053.

[2] P. D'Ambra and P. S. Vassilevski, *Adaptive AMG with coarsening based on compatible weighted matching,* Available as Lawrence Livermore National Laboratory Technical Report LLNL-TR-613612, January 28, 2013.

[3] I. S. Duff, J. Koster, *On algorithms for permuting large entries to the diagonal of a sparse matrix*, SIAM Journal on Matrix Analysis and Applications **22**(4)(2001), pp. 973–996.

[4] HSL(2011). *A collection of Fortran codes for large scale scientific computation. http://www.hsl.rl.ac.uk.*

[5] M.E.J. Newman, *Networks. An Introduction,* Oxford University Press, 2010.

[6] J. Yang and J. Leskovec. *Defining and Evaluating Network Communities based on Ground-truth* ICDM, 2012.

A Parallel Geometric Multigrid Solver for Density Driven Flow

Andreas Vogel
(joint work with Sebastian Reiter, Gabriel Wittum)

Multigrid methods for the solution of large sparse matrices arising from grid-based discretizations of partial differential equations are well known for their optimal complexity, i.e., the computation effort only increases linearly with the problem size. This makes them a promising algorithm when focusing on the weak scaling properties of such a matrix solver. However, while reducing the problem size within a multigrid algorithm on coarser grid levels is its strength, this gives rise to a potential performance bottleneck when parallelization is taken into account. Indeed, on coarser grid levels the inner to boundary ratio of the grid parts assigned to a process become unpleasant and a parallel smoother on those coarse levels will suffer from the fact that mostly communication at the boundary takes place and only little computation on the inner part is performed. In order to overcome this bottleneck we present an algorithm that avoids this situation by gathering coarser levels to fewer processors leaving the remaining processors idle. To this end we introduce vertical interface connections that allow this gathering process and adapt the transfer operators of the multigrid algorithm to respect these interfaces. Arriving at a single process on the coarsest level a serial base solver, e.g., LU factorization, can be used. We show that this approach leads to nice weak scaling behavior for an exemplary application: Discretizing a pde system for density driven flow using a vertex-centered finite volume scheme and implicit Euler time stepping we analyze the efficiency of the geometric multigrid solver in the first Newton linearization of the first time step. It turns out that up to 130,000 processors the weak scaling efficiency is still above 80%.

Adaptive inexact Newton methods with a posteriori stopping criteria for nonlinear diffusion PDEs

Martin Vohralík
(joint work with Alexandre Ern)

We consider nonlinear algebraic systems resulting from numerical discretizations of nonlinear partial differential equations. To solve these systems, some iterative nonlinear solver, and, on each step of this solver, some iterative linear solver are used. We derive adaptive stopping criteria for both iterative solvers. Our criteria are based on an a posteriori error estimate which distinguishes the different error components, namely the discretization error, the linearization error, and the algebraic error. We stop the iterations whenever the corresponding error does no longer affect the overall error significantly. Our estimates also yield a guaranteed upper bound on the overall error at each step of the nonlinear and linear solvers. We prove the (local) efficiency and robustness of the estimates with respect to the size of the nonlinearity owing, in particular, to the error measure involving the dual norm of the residual. Our developments hinge on equilibrated flux reconstructions

that we detail in the linear case. We present a general abstract framework yielding a guaranteed upper bound and local efficiency with polynomial degree robustness and show how to apply this framework to various discretization schemes like finite elements, nonconforming finite elements, discontinuous Galerkin, or mixed finite elements. Numerical experiments illustrate the tight overall error control and important computational savings achieved in our approach. Details can be found in [1, 2, 3].

References

[1] A. Ern and M. Vohralík, *Adaptive inexact Newton methods: a posteriori error control and speedup of calculations*, SIAM News **46** (2013), 1,4.
[2] A. Ern and M. Vohralík, *Adaptive inexact Newton methods with a posteriori stopping criteria for nonlinear diffusion PDEs*, SIAM J. Sci. Comput. **35** (2013), A1761–A1791.
[3] A. Ern and M. Vohralík, *Polynomial-degree-robust a posteriori estimates in a unified setting for conforming, nonconforming, discontinuous Galerkin, and mixed discretizations*, HAL Preprint 00921583 (2013).

Fast Multilevel Methods for the Helmholtz Equation

Kees Vuik

Many wave phenomena are well described by the wave equation. When the considered wave has a fixed frequency the wave equation is mostly re-written in the frequency domain which results in the Helmholtz equation. It is also possible to approximate the time domain solution with a summation of solutions for several frequencies. Applications are the propagation of sound, sonar, seismics, and many more. We take as an example the search for oil and gas using seismics. It is well known that an increase of the frequency leads to a higher resolution, so more details of the underground become visible.

The Helmholtz equation in its most simple form is a combination of the symmetric positive Poisson operator and a negative constant, the so-called wavenumber, multiplied with the identity operator. In order to find the solution in a complicated domain a discretization has to be done. There are two characteristic properties of the discretized system:

- the product of the wavenumber and the step size should be smaller than a given constant
- if the wavenumber increases the operator has more and more negative eigenvalues.

If damping is involved the operator also has a part with an imaginary value. The resulting matrix is however not Hermitian.

In the years around 2000 no good iterative solvers are known. The standard approaches: Krylov and multi-grid break down if the wavenumber increases.

Around 2005 a new preconditioner based on the shifted Laplace preconditioner was proposed, which leads to a class of fast and robust Helmholtz solvers. It appears that the amount of work increases linearly with the wavenumber. At this moment, this is the method of choice to solve the Helmholtz equation. Various papers have appeared to analyze the good convergence behavior.

Recently a multi level Krylov solver has been proposed that seems to be scalable, which means that the number of iterations is also independent of the wavenumber (see Figure 1). An analysis of this method is given and recent results for industrial problems are presented.

Finally implementation of the method on modern High Performance Computing platforms will be discussed.

References

[EOV06] Y.A. Erlangga and C.W. Oosterlee and C. Vuik A Novel Multigrid Based Preconditioner For Heterogeneous Helmholtz Problems SIAM J. Sci. Comput.,27, pp. 1471-1492, 2006 http://ta.twi.tudelft.nl/nw/users/vuik/papers/Erl06OV.pdf

[EN08] Y.A. Erlangga and R. Nabben On a multilevel Krylov method for the Helmholtz equation preconditioned by shifted Laplacian ETNA, 31, pp. 403-424, 2008

[KOV11] H. Knibbe and C.W. Oosterlee and C. Vuik GPU implementation of a Helmholtz Krylov solver preconditioned by a shifted Laplace multigrid method Journal of Computational and Applied Mathematics, 236, pp. 281-293, 2011 http://ta.twi.tudelft.nl/nw/users/vuik/papers/Kni11OV.pdf

[SHV13] A.H. Sheikh and D. Lahaye and C. Vuik On the convergence of shifted Laplace preconditioner combined with multilevel deflation Numerical Linear Algebra with Applications, 20, pp. 645-662, 2013 http://ta.twi.tudelft.nl/nw/users/vuik/papers/She13LV.pdf

BDDC deluxe domain decomposition algorithms

Olof B. Widlund
(joint work with Clark R. Dohrmann)

The BDDC algorithms, first developed by Clark Dohrmann, have proven to be very successful domain decomposition algorithms for a variety of elliptic problems. For any particular application, the success of such an algorithm depends on the choice of a set of primal constraints and on the choice of an averaging operator, which is used to restore the continuity of certain intermediate vectors in each iteration.

In the deluxe version, a new averaging procedure is used; it was first developed in joint work with Dohrmann on H(curl) problems in order to deal with the two material parameters of such problems. The algorithm will be described and theory will be outlined. Several successful applications will be discussed including problems posed in H(div), Reissner-Mindlin plates, and elliptic problems and isogeometric analysis.

Efficient solution methods for space-time discontinuous Petrov-Galerkin discretizations for the wave equation

Christian Wieners

We consider discretizations of linear wave equations (acoustic, elastic or electromagnetic waves) using higher order discontinuous Galerkin elements with upwind flux in space. In the first step, we compare the performance of explicit and implicit Runge-Kutta methods with polynomial and rational Krylov methods for uniform meshes and fixed time steps [1]. Then, this is extended to a space-time adaptive variational Petrov-Galerkin space-time discretization with continuous Galerkin elements in time. Finally, we show that the DPG method can be also transferred to this problem class by applying the setting presented in [3].

The wave equation. We consider linear hyperbolic systems $M\partial_t \mathbf{u}(t)+A\mathbf{u}(t) = 0$ for $t \in [0,T]$ defined on the domain $V \subset \mathrm{L}_2(\Omega)^J$ of the operator A; here, $\Omega \subset \mathbb{R}^D$ is a bounded Lipschitz domain. Our main example is the Maxwell system, where $M(\mathbf{H},\mathbf{E}) = (\mu\mathbf{H}, \varepsilon\mathbf{E})$ and $A(\mathbf{H},\mathbf{E}) = (\operatorname{curl}\mathbf{E}, -\operatorname{curl}\mathbf{H})$. With suitable boundary conditions in V, the operator $-M^{-1}A$ generates a semigroup in V. The operator A is approximated by A_h in a dG space V_h using full upwind flux.

Time integrators. The solution $\mathbf{u}_h(t) = \exp(-tM_h^{-1}A_h)\mathbf{u}_h(0)$ can be approximated by explicit Runge-Kutta methods, provided the time step is small enough. Implicit Gauss collocation methods are unconditionally stable, but the solution of linear systems are required.

Alternatively, the discrete matrix exponential can be approximated by Krylov methods: for an M-orthogonal basis $Q_m = (\mathbf{q}_1, \dots, \mathbf{q}_m)$ of the Krylov subspace we obtain $\exp(-\triangle t M_h^{-1}A_h) \approx Q_m \exp(-\triangle t Q_m^\top A_h Q_m) Q_m^\top$; this can be evaluated efficiently for $m \ll \dim V_h$.

Numerical experiments for the 2D Maxwell system show that—due to the CFL limitation—explicit methods are not competitive, if suitable solution methods for implicit time integrators are available [1, Sect. 6].

Adaptive Petrov-Galerkin space-time discretizations. In order to include adaptivity in space and time simultaneously, we consider a variational Petrov-Galerkin space-time discretization with continuous Galerkin elements in time for the linear operator $L = M\partial_t + A$ with domain $U \subset H = \mathrm{L}_2(\Omega \times (0,T))$. The inf-sup condition can be verified in this setting, so that standard theory applies. Adaptivity is controlled by a goal-oriented dual weighed error estimator: for the solution $\mathbf{u} \in U$ of $L\mathbf{u} = \mathbf{f}$ and a given linear error functional $\mathcal{J}(\cdot)$, the error $\mathbf{e} = \mathbf{u} - \mathbf{u}_h$ satisfies

$$\mathcal{J}(\mathbf{e}) = (L\mathbf{e}, \mathbf{u}^* - \mathbf{u}_h^*)_{\Omega\times(0,T)} = (\mathbf{f} - L\mathbf{u}_h, \mathbf{u}^* - \mathbf{u}_h^*)_{\Omega\times(0,T)}\,,$$

where $\mathbf{u}^*$ is the dual solution determined by

$$(L\mathbf{v}, \mathbf{u}^*)_{\Omega\times(0,T)} = \mathcal{J}(\mathbf{v})\,, \qquad \forall \mathbf{v}\,.$$

The unknown exact dual solution $\mathbf{u}^*$ is approximated from the computed solution $\mathbf{u}_h^*$ via a polynomial recovery of higher order in space and time.

Numerical examples, again for the 2D Maxwell case, demonstrate the efficiency of the error estimation, the adaptive strategy, and the parallel solution process. These results are a joint work with Stefan Findeisen and Willy Dörfler. All computations are realized within the parallel finite element software system M++ [2].

A discontinuous Petrov-Galerkin space-time discretization. This method is based on a decomposition $\bigcup_{\tau\in\mathcal{T}}\bar\tau$ into space-time cells $\tau = K_\tau \times (t_\tau^{\min}, t_\tau^{\max})$. Let $U_\tau = U|_\tau$ be the restriction to τ, and let $\hat{\mathbf{u}}_\tau = \gamma_\tau \mathbf{u}_\tau$ be the trace

$$\hat{\mathbf{u}}_\tau(\mathbf{x},t) = \begin{cases} -\mathbf{u}_\tau(\mathbf{x},t_{\min}) & \mathbf{x}\in K\,, \\ \mathbf{u}_\tau(\mathbf{x},t_{\max}) & \mathbf{x}\in K\,, \\ \gamma_K \mathbf{u}_\tau(\mathbf{x},t) & \mathbf{x}\in\partial K \text{ and } t\in(t_{\min},t_{\max})\,. \end{cases}$$

Let $\gamma_\tau^{\mathrm{ad}}$ be the adjoint trace mapping and U^{ad} the adjoint space with $\mathbf{u}(T)=0$. Integration by parts yields $L^{\mathrm{ad}} = -L$ and

$$(L\mathbf{u},\mathbf{v})_\tau = (\mathbf{u}, L^{\mathrm{ad}}\mathbf{v})_\tau + \langle \gamma_\tau \mathbf{u}, \gamma_\tau^{\mathrm{ad}}\mathbf{v}\rangle\,, \qquad \mathbf{u}\in U_\tau\,,\ \mathbf{v}\in U_\tau^{\mathrm{ad}}\,.$$

Define $U_{\mathcal{T}}^{\mathrm{ad}} = \prod U_\tau^{\mathrm{ad}}$, $\gamma_{\mathcal{T}} = (\gamma_\tau)$, the trace space $\hat U = \gamma_{\mathcal{T}}(U) \subset \prod U_\tau/\mathcal{N}(\gamma_\tau)$, and

$$b\colon \hat U \times H \times U_{\mathcal{T}}^{\mathrm{ad}} \longrightarrow \mathbb{R}\,, \qquad b(\hat{\mathbf{u}},\mathbf{u},\mathbf{v}) = \sum\nolimits_\tau (\mathbf{u}, L^{\mathrm{ad}}\mathbf{v})_\tau + \langle \hat{\mathbf{u}}_\tau, \gamma_\tau^{\mathrm{ad}}\mathbf{v}\rangle\,.$$

For example, we have for acoustic waves with $L(\mathbf{q},p) = (\partial_t\mathbf{q} + \nabla p, \rho\partial_t p + \nabla\cdot\mathbf{q})$

$$\begin{aligned}(\partial_t\mathbf{q}+\nabla p, \tilde{\mathbf{q}})_\tau + (\rho\partial_t p + \nabla\cdot\mathbf{q}, \tilde q)_\tau &= -(\mathbf{q}, \partial_t\tilde{\mathbf{q}} + \nabla\tilde p)_\tau - (p, \rho\partial_t\tilde p + \nabla\cdot\tilde{\mathbf{q}})_\tau \\ &+\big(\mathbf{q}(t_{\max}), \tilde{\mathbf{q}}(t_{\max})\big)_K - \big(\mathbf{q}(t_{\min}), \tilde{\mathbf{q}}(t_{\min})\big)_K + (\nabla\mathbf{q}\cdot\mathbf{n}, \tilde p)_{\partial K\times(t_{\min},t_{\max})} \\ &+\big(p(t_{\max}), \tilde p(t_{\max})\big)_K - \big(p(t_{\min}), p(t_{\min})\big)_K + (p, \nabla\tilde{\mathbf{q}}\cdot\mathbf{n})_{\partial K\times(t_{\min},t_{\max})}\,.\end{aligned}$$

Let $\hat U_h \subset \hat U_h$ be a discrete trace space, let $H_h = \prod H_{\tau,h} \subset H$ be a discrete ansatz space, and let $U_{\mathcal{T},h}^{\mathrm{ad}} = \prod U_{\tau,h}^{\mathrm{ad}} \subset U_{\mathcal{T}}^{\mathrm{ad}}$ be a discrete broken test space such that

$$\sup_{\mathbf{v}_{\tau,h}\in U_{\tau,h}^{\mathrm{ad}}} \frac{b(\hat{\mathbf{u}}_h, \mathbf{u}_{\tau,h}, \mathbf{v}_{\tau,h})}{\|\mathbf{v}_{\tau,h}\|_{U_\tau^{\mathrm{ad}}}} \geq \beta_0 \sup_{\mathbf{v}_\tau\in U_\tau^{\mathrm{ad}}} \frac{b(\hat{\mathbf{u}}_h, \mathbf{u}_{\tau,h}, \mathbf{v}_\tau)}{\|\mathbf{v}_\tau\|_{U_\tau^{\mathrm{ad}}}}$$

for all $(\hat{\mathbf{u}}_h, \mathbf{u}_{\tau,h}) \in \hat U_h \times H_{\tau,h}$ and all $\tau\in\mathcal{T}$. We define the optimal test space

$$U_{\mathcal{T},h}^{\mathrm{opt}} = \Big\{ \mathbf{v}_h \in U_{\mathcal{T},h}^{\mathrm{ad}}\colon (\hat{\mathbf{u}}_h,\mathbf{u}_h)\in\hat U_h\times H_h \text{ exists such that } (\mathbf{v}_h,\mathbf{w}_h)_{U_{\mathcal{T}}^{\mathrm{ad}}} = b(\hat{\mathbf{u}}_h,\mathbf{u}_h,\mathbf{w}_h) \text{ for all } \mathbf{w}_h\in U_{\mathcal{T},h}^{\mathrm{ad}} \Big\}\,.$$

Then, for $\mathbf{f}\in \mathrm{L}_2(Q)^J$, a unique solution $(\hat{\mathbf{u}}_h,\mathbf{u}_h)\in\hat U_h\times H_h$ exists solving

$$b(\hat{\mathbf{u}}_h,\mathbf{u}_h,\mathbf{v}_h) = (\mathbf{f},\mathbf{v}_h)_Q\,, \qquad \mathbf{v}_h\in U_{\mathcal{T},h}^{\mathrm{opt}}\,,$$

and for the discretization error holds (cf. [3, Th. 6])

$$\|(\hat{\mathbf{u}},\mathbf{u}) - (\hat{\mathbf{u}}_h,\mathbf{u}_h)\|_{\hat U\times H} \leq C \inf_{(\hat{\mathbf{w}}_h,\mathbf{w}_h)\in\hat U_h\times H_h} \|(\hat{\mathbf{u}},\mathbf{u}) - (\hat{\mathbf{w}}_h,\mathbf{w}_h)\|_{\hat U\times H}\,.$$

References

[1] M. Hochbruck, T. Pažur, A. Schulz, E. Thawinan, and C. Wieners. Efficient time integration for discontinuous Galerkin approximations of linear wave equations. Technical report, Karlsruhe Institute of Technology, 2013. http://na.math.kit.edu/download/papers/TI-DG-Wave.pdf.

[2] C. Wieners, *A geometric data structure for parallel finite elements and the application to multigrid methods with block smoothing*, Comput. Vis. Sci., 13 (2010), pp. 161–175.

[3] C. Wieners and B. Wohlmuth. Robust operator estimates and the application to substructuring methods for first-order systems. to appear in M2AN http://www.math.kit.edu/user/ wieners/WienersWohlmuthRevised13.pdf.

Dual and Hybrid Hierarchical Grids for Fast Geophysical Flow Solvers

Barbara Wohlmuth

(joint work with B. Gmeiner, H. Stengel, U. Rüde, C. Waluga)

Efficient and fast iterative solution methods for Stokes-type systems are a crucial ingredient for coupled geophysical solvers, e.g., in mantle-convection- or ice-sheet modeling. Recent publications emphasize the importance of the interplay between the physical model and suitable discretization concepts, involving matching iterative solvers and parallelization strategies. In an effort to tackle next-generation challenges and computer architectures, we propose a strategy for geophysical flow solvers which is based on a careful combination of both existing and specifically designed numerical techniques. To simplify the exposition, we neglect compressibility effects and nonlinearities in the material parameters, and consider the following dimensionless model in a polyhedral domain $\Omega \subset \mathbb{R}^d$, $d = 2, 3$. The problem is defined by the conservation of momentum/mass and energy via

$$\operatorname{div}(\mu(\nabla \mathbf{u} + \nabla \mathbf{u}^\top) - p\mathbf{I}) = \mathrm{Ra} T \widehat{\mathbf{r}}, \qquad \operatorname{div} \mathbf{u} = 0, \tag{1}$$

$$\partial_t T + \mathbf{u} \cdot \nabla T = \mathrm{Pe}^{-1} \Delta T, \tag{2}$$

respectively, where $\mathbf{u}$ denotes the velocity, p the pressure, T the temperature, $0 < \mu \in L^\infty(\Omega)$ the viscosity, and $\widehat{\mathbf{r}}$ the unit vector in the direction of gravity. The Rayleigh number Ra and the Peclet number Pe determine the presence, vigourousity, and strength of convection in the system and are typically assumed to take extreme values, which may require appropriate stabilization of all three equations in discretized form. The model is complemented by suitable initial and boundary conditions. That being said, the knowledge of these conditions is in practice severely limited. One motivation of having extreme-scale iterative solvers is the inversion of models of the above type alongside with the quantification of parameter and model uncertainties, requiring a large number of solves.

Semi-discretization in time: The state-of-the-art in geophysical flow computations are explicit time-stepping methods such as Runge–Kutta type schemes. This, on one hand, decouples the problem into the explicit evolution of the temperature and an instantaneous constraint governed by the Stokes problem resulting in an implicit problem. On the other hand, the time-stepping is CFL-limited, which

drives the number of auxiliary problems (1) to be solved in high-resolution computations. Therefore, it is not only important that we reach extreme scalability, but also deliver a small time-to-solution for practical applications.

Treatment of the Stokes part (1)**:** Despite the vast literature dealing with the discretization and solution of the Stokes problem, the majority of the published work considers only the two-dimensional in practical realizations. Moreover, generic constants appearing in the analysis of discretization schemes and solvers are often left unspecified, thereby only giving information about their asymptotic optimality but not of the actual performance. The lack of fair comparisons between different three-dimensional velocity-pressure pairings with respect to quality of the approximation vs. number of degrees of freedom and sparsity of the resulting system makes it difficult to rule out potential candidates and decide upon a promising choice. While continuous pressure schemes do not straightforwardly yield desirable physical properties like strong or at least local mass-conservation, most low order discontinuous pressure schemes on tetrahedral meshes have serious stability issues (e.g., $P_2 - P_0$, or Crouzeix-Raviart for non-isoviscous flow). For higher orders, we may require bubble-enrichments or barycentrically refined meshes, both of which do not lend themselves to iterative solvers based on hierarchic decompositions.

Due to the circumstances mentioned above, we base our solver prototype on equal-order linear interpolations for the velocity and the pressure which allow the efficient implementation of matrix-free multigrid. Let us remark that in a hierarchical setting, linear finite elements for the Stokes problem have the additional advantage that the pressure converges with $\mathcal{O}(h^{3/2})$ due to a super-closeness property, which adds to the ratio of accuracy per degrees of freedom desired in extreme-scale high-performance computations.

Since these discretizations are neither locally mass-conservative nor stable, we propose a combination of stabilization and local postprocessing to circumvent these problems. For stability reasons, we place the pressure degrees of freedom on a coarser level which yields a stable discretization and allows the later extension to quadratic Taylor–Hood elements. For the sake of local mass-conservation, we use a novel method to compute face-wise constant mass-flux-corrections on dual cells [2]. By solving small local systems, we can moreover lift the velocity to strongly divergence-free approximations on the primal mesh.

The linear systems resulting from this discretization of the Stokes problem (1) are solved in a nested iterative fashion, using a preconditioned CG on the pressure Schur complement for the outer iterations, and geometric multigrid for the (approximate) inversion of the viscous operator. While this pressure-correction scheme may not deliver optimal convergence rates, we can achieve good times to solution using this approach, especially in unsteady computations. The thorough comparison with specialized saddle-point smoothers (e.g., Vanka or DGS) and the development of robust multigrid techniques for the varying viscosity case are subjects of ongoing research.

Treatment of the energy part (2)**:** Given a postprocessed velocity field, we can achieve local energy conservation by using suitable discretizations for the

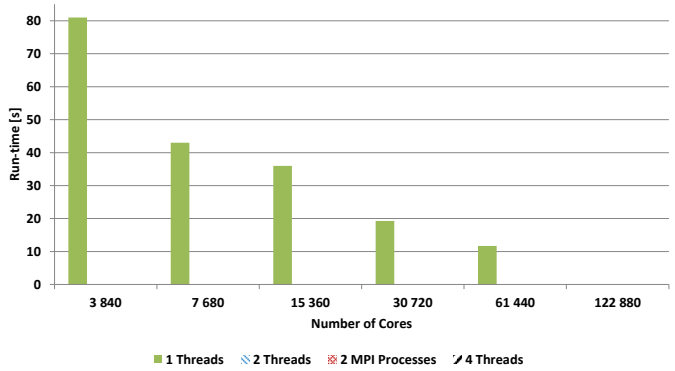

FIGURE 1. Strong scaling of the Stokes-solver on JUQUEEN.

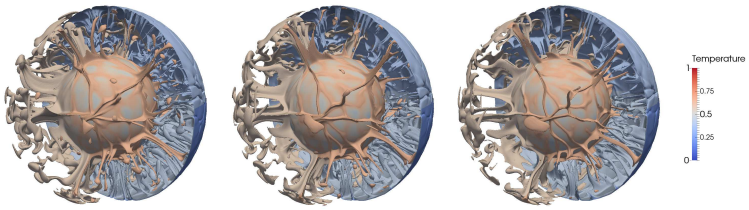

FIGURE 2. Temperature iso-surfaces inside a spherical shell.

energy part of the problem. Since nodal data-structures are already used for the Stokes-part, a natural choice here is the use of vertex-centered finite-volume schemes on dual meshes. The coupling between postprocessed equal-order discretizations and finite volumes was demonstrated in 2D computations in [2]. The conservative coupling in 3D follows by the same techniques, and a practical realization is work in progress. For the time being, we use a streamline-upwinding Petrov–Galerkin approach which is not locally conservative and requires parameter tuning. The reason for this is that such a finite element based scheme can be easily realized using the same data-structures as the vertex-centered finite volumes and is therefore suitable for first performance studies.

Computational results: Our highly optimized implementation in the hierarchical hybrid grid framework (HHG) demonstrated the feasibility of solving Stokes-systems with $\mathcal{O}(10^{12})$ unknowns in the range of minutes [1]. The solver showed excellent scalability on state-of-the-art peta-scale architectures; cf. e.g. Fig. 1. Moreover, we demonstrated the performance of the solver for Boussinesq-type flow in a spherical shell geometry involving $\mathcal{O}(10^{10})$ unknowns and $10,000$ time-steps; cf. Fig. 2.

References

[1] B. Gmeiner, U. Rüde, H. Stengel, C. Waluga, and B. Wohlmuth, *Performance and Scalability of Hierarchical Hybrid Multigrid Solvers for Stokes Systems*, submitted (2013).
[2] B. Gmeiner, C. Waluga, and B. Wohlmuth. *Local mass-conservation for stabilized linear finite element interpolations of incompressible flow*, submitted (2014)

Reporter: Markus Knodel

Participants

Dr. Markus Bachmayr
Institut f. Geometrie & Prakt. Mathematik
RWTH Aachen
Templergraben 55
52056 Aachen
GERMANY

Prof. Dr. Randolph E. Bank
Department of Mathematics
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0112
UNITED STATES

Prof. Dr. Steffen Börm
Christian-Albrechts-Universität für Informatik
Lehrstuhl Scientific Computing
Christian-Albrecht-Platz 4
24118 Kiel
GERMANY

Prof. Dr. Dietrich Braess
Fakultät für Mathematik
Ruhr-Universität Bochum
44780 Bochum
GERMANY

Markus Breit
G-CSC
Wolfgang-Goethe-Universität Frankfurt/M.
Im Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Prof. Dr. Susanne C. Brenner
Department of Mathematics
Louisiana State University
Baton Rouge LA 70803-4918
UNITED STATES

Prof. Dr. Dr.h.c. Peter Deuflhard
Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB)
Takustr. 7
14195 Berlin
GERMANY

Prof. Dr. Craig C. Douglas
Department of Mathematics
Institute of Scientific Computation
University of Wyoming
Box 3036 University Station
Laramie, WY 82071
UNITED STATES

Dr. Robert D. Falgout
Center of Applied Scientific Computing
Lawrence Livermore National Laboratory
P.O.Box 808, L-561
Livermore, CA 94551
UNITED STATES

Dr. Björn Gmeiner
Zentrum Mathematik
TU München
Boltzmannstr. 3
85748 Garching b. München
GERMANY

Prof. Dr. Lars Grasedyck
Institut f. Geometrie & Prakt. Mathematik
RWTH Aachen
52056 Aachen
GERMANY

Dr. Alfio Grillo
Politecnico di Torino
DISMA - Dipartimento di Scienze
Matematiche "G. L. Lagrange"
Corso Duca degli Abruzzi, 24
10129 Torino
ITALY

Univ.-Prof. Dr. Gundolf Haase
Institut für Mathematik und
wissenschaftliches Rechnen
Karl-Franzens-Universität Graz
Heinrichstr. 36
8010 Graz
AUSTRIA

Prof. Dr. Dr.h.c. Wolfgang Hackbusch
Max-Planck-Institut für Mathematik
in den Naturwissenschaften
Inselstr. 22 - 26
04103 Leipzig
GERMANY

Michael Hoffer
G-CSC
Wolfgang-Goethe-Universität
Frankfurt/M.
Im Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Susanne Höllbacher
Goethe-Center for Scientific Computing
Goethe Universität Frankfurt
Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Prof. Dr. Ronald H. W. Hoppe
Lehrstuhl f. Angewandte Mathematik I
Universität Augsburg
Universitätsstr. 14
86159 Augsburg
GERMANY

Prof. Dr. Klaus Johannsen
Uni Computing, Director
Uni Research AS
Thormöhlensgt. 55
5008 Bergen
NORWAY

Prof. Dr. David Elliot Keyes
Appl. Physics and Appl. Mathematics
Columbia University
200 S.W. Mudd Bldg., MC 4701
500 W. 120th Street
New York, NY 10027
UNITED STATES

Dr. Boris N. Khoromskij
Max-Planck-Institut für Mathematik
in den Naturwissenschaften
Inselstr. 22 - 26
04103 Leipzig
GERMANY

Dr. Markus Knodel
G-CSC
Wolfgang-Goethe-Universität
Frankfurt/M.
Im Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Prof. Dr. Ralf Kornhuber
Institut für Mathematik
Freie Universität Berlin
Arnimallee 6
14195 Berlin
GERMANY

Prof. Dr. Rolf Krause
Universita della Svizzera Italiana
Facolta di Scienze Informatiche
Via Giuseppe Buffi 13
6904 Lugano
SWITZERLAND

Prof. Dr. Ulrich Langer
Institut für Numerische Mathematik
Johannes Kepler Universität Linz
Altenbergerstr. 69
4040 Linz
AUSTRIA

Prof. Dr. Sabine Le Borne
TU Hamburg-Harburg
Institut für Mathematik (E-10)
Schwarzenbergstrasse 95
21073 Hamburg
GERMANY

Babett Lemke
G-CSC
Wolfgang-Goethe-Universität
Frankfurt/M.
Im Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Christian Löbbert
Institut f. Geometrie & prakt.
Mathematik
RWTH Aachen
Templergraben 55
52056 Aachen
GERMANY

Prof. Dr. Volker Mehrmann
Institut für Mathematik
Technische Universität Berlin
Sekr. MA 4-5
Strasse des 17. Juni 136
10623 Berlin
GERMANY

Prof. Dr. Arnd Meyer
Fakultät für Mathematik
TU Chemnitz
09107 Chemnitz
GERMANY

Prof. Dr. Hans D. Mittelmann
School of Mathematical & Statistical
Sciences
Arizona State University
Box 871804
Tempe, AZ 85287-1804
UNITED STATES

Dr. Arne Nägel
G-CSC
Goethe-Universität Frankfurt/M.
Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Prof. Dr. Jeffrey S. Ovall
Fariborz Maseeh Department
of Mathematics and Statistics
Portland State University
315 Neuberger Hall
Portland, OR 97207-0751
UNITED STATES

Prof. Dr. Gillian Queisser
Goethe-Center for Scientific Computing
Goethe Universität Frankfurt
Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Dr. Christoph Reisinger
Mathematical Institute
Oxford University
24-29 St. Giles
Oxford OX1 3LB
UNITED KINGDOM

Sebastian Reiter
G-CSC
Wolfgang-Goethe-Universität
Frankfurt/M.
Im Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Prof. Dr. Arnold Reusken
Institut f. Geometrie & Prakt. Mathematik
RWTH Aachen
Templergraben 55
52061 Aachen
GERMANY

Prof. Dr. Stefan A. Sauter
Institut für Mathematik
Universität Zürich
Winterthurerstr. 190
8057 Zürich
SWITZERLAND

Prof. Dr. Volker Schulz
Fachbereich IV - Mathematik, Numerik, Optimierung & part. Differentialgleichungen
Universität Trier
54286 Trier
GERMANY

Prof. Dr. Ridgway Scott
Department of Mathematics
The University of Chicago
257F Ryerson Hall
Chicago, IL 60637
UNITED STATES

Prof. Dr. Rob Stevenson
Korteweg-de Vries Instituut
Universiteit van Amsterdam
Postbus 94248
1090 GE Amsterdam
NETHERLANDS

Prof. Dr. Stefan Vandewalle
Department of Computer Science
KU Leuven
Celestijnenlaan 200A
3001 Leuven
BELGIUM

Prof. Dr. Wim Vanroose
Department Wiskunde-Informatica
University of Antwerpen
Middelheimlaan 1
2020 Antwerpen
BELGIUM

Prof. Dr. Panayot S. Vassilevski
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
P.O.Box 808, L-560
Livermore CA 94550
UNITED STATES

Andreas Vogel
G-CSC
Wolfgang-Goethe-Universität Frankfurt/M.
Im Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Dr. Martin Vohralik
INRIA Rocquencourt
Domaine de Voluceau
B. P. 105
78153 Le Chesnay Cedex
FRANCE

Prof. Dr. Kees Vuik
Delft Institute of Applied Mathematics
Delft University of Technology
Mekelweg 4
2628 Delft CD
NETHERLANDS

Prof. Dr. Olof B. Widlund
Courant Institute of Math. Sciences
New York University
251, Mercer Street
New York, NY 10012-1110
UNITED STATES

Prof. Dr. Christian Wieners
Karlsruher Institut f. Technologie (KIT)
Inst. f. Angew. & Numerische Mathematik
76131 Karlsruhe
GERMANY

Prof. Dr. Harry Yserentant
Institut für Mathematik
Technische Universität Berlin
Straße des 17. Juni 136
10623 Berlin
GERMANY

Prof. Dr. Gabriel Wittum
G-CSC
Wolfgang-Goethe-Universität Frankfurt/M.
Im Kettenhofweg 139
60325 Frankfurt am Main
GERMANY

Prof. Dr. Christoph Zenger
Institut für Informatik
TU München
Boltzmannstr. 3
85748 Garching
GERMANY

Prof. Dr. Barbara Wohlmuth
Zentrum Mathematik
TU München
Boltzmannstr. 3
85748 Garching b. München
GERMANY