

An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earth’s Mantle

Johann Rudi*, A. Cristiano I. Malossi†, Tobin Isaac*, Georg Stadler‡, Michael Gurnis‡, Peter W. J. Staar†, Yves Ineichen†, Costas Bekas†, Alessandro Curioni†, Omar Ghattas*¶

*Institute for Computational Engineering and Sciences, The University of Texas at Austin, USA

†Foundations of Cognitive Solutions, IBM Research – Zurich, Switzerland

‡Courant Institute of Mathematical Sciences, New York University, USA

‡Seismological Laboratory, California Institute of Technology, USA

¶Jackson School of Geosciences and Dept. of Mechanical Engineering, The University of Texas at Austin, USA

Abstract—Mantle convection is the fundamental physical process within earth’s interior responsible for the thermal and geological evolution of the planet, including plate tectonics. The mantle is modeled as a viscous, incompressible, non-Newtonian fluid. The wide range of spatial scales, extreme variability and anisotropy in material properties, and severely nonlinear rheology have made global mantle convection modeling with realistic parameters prohibitive. Here we present a new implicit solver that exhibits optimal algorithmic performance and is capable of extreme scaling for hard PDE problems, such as mantle convection. To maximize accuracy and minimize runtime, the solver incorporates a number of advances, including aggressive multi-octree adaptivity, mixed continuous-discontinuous discretization, arbitrarily-high-order accuracy, hybrid spectral/geometric/algebraic multigrid, and novel Schur-complement preconditioning. These features present enormous challenges for extreme scalability. We demonstrate that—contrary to conventional wisdom—algorithmically optimal implicit solvers can be designed that scale out to 1.5 million cores for severely nonlinear, ill-conditioned, heterogeneous, and anisotropic PDEs.

Submission Category: Scalability

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

SC '15, November 15 - 20, 2015, Austin, TX, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3723-6/15/11...\$15.00

DOI: <http://dx.doi.org/10.1145/2807591.2807675>

I. EARTH’S MANTLE CONVECTION

Earth is a dynamic system in which mantle convection drives plate tectonics and continental drift and, in turn, controls much activity ranging from the occurrence of earthquakes and volcanoes to mountain building and long-term sea level change. Despite its central role in solid earth dynamics, we have enormous first-order gaps in our knowledge of mantle convection, with questions that are as basic as what are the principal driving and resisting forces on plate tectonics to what is the energy balance of the planet as a whole. Indeed, understanding mantle convection has been designated one of the “*10 Grand Research Questions in Earth Sciences*” in a recent National Academies report [1]. We seek to address such fundamental questions as: (i) What are the main drivers of plate motion—negative buoyancy forces or convective shear traction? (ii) What is the key process governing the occurrence of great earthquakes—the material properties between the plates or the tectonic stress?

Addressing these questions requires global models of earth’s mantle convection and associated plate tectonics, with realistic parameters and high resolutions down to faulted plate boundaries. Historically, modeling at this scale has been out of the question due to the enormous computational complexity associated with numerical solution of the underlying mantle flow equations. However, with the advent of multi-petaflops supercomputers as well as significant advances in seismic tomography and space geodesy placing key observational constraints on mantle convection, we now have the opportunity to address these fundamental questions.

Instantaneous flow of the mantle is modeled by the




Figure 1: Cross section through a subducting slab of the Pacific plate showing effective mantle viscosity (colors) in a simulation of nonlinear mantle flow. Viscosity of plates reaches μ_{\max} (dark blue), except at the trench due to plastic yielding. In the thin plate boundary region, viscosity drops to μ_{\min} (dark red) creating a contrast of 10^6 . Strain rate weakening reduces the viscosity underneath the plates and, combined with plastic yielding, is the reason for earth’s highly nonlinear rheology.

nonlinear incompressible Stokes equations:

$$-\nabla \cdot \left[\mu(T, \mathbf{u}) (\nabla \mathbf{u} + \nabla \mathbf{u}^\top) \right] + \nabla p = \mathbf{f}(T), \quad (1a)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (1b)$$

where \mathbf{u} , T , and p are the velocity, temperature, and pressure fields, respectively; \mathbf{f} is the temperature-dependent forcing derived from the Boussinesq approximation; and the temperature- and velocity-dependent effective viscosity μ is characterized by the constitutive law

$$\mu(T, \mathbf{u}) = \mu_{\min} + \quad (1c)$$

$$\min \left(\frac{\tau_{\text{yield}}}{2\dot{\epsilon}_{II}(\mathbf{u})}, w \min \left(\mu_{\max}, a(T) \dot{\epsilon}_{II}(\mathbf{u})^{1/n-1} \right) \right).$$

The effective viscosity depends on a power of the square root of the second invariant of the strain rate tensor $\dot{\epsilon}_{II} := \frac{1}{2}(\dot{\epsilon} : \dot{\epsilon})^{1/2}$, where “:” represents the inner product of second order tensors and $\dot{\epsilon} := \frac{1}{2}(\nabla \mathbf{u} + \nabla \mathbf{u}^\top)$. The viscosity decays exponentially with temperature via the Arrhenius relationship, symbolized by $a(T)$. The constitutive relation incorporates plastic yielding with yield stress τ_{yield} , lower/upper bounds on viscosity μ_{\min} and μ_{\max} , and a decoupling factor $w(\mathbf{x})$ to model plate boundaries. For whole earth mantle flow models, (1) are augmented with free-slip conditions (i.e., no tangential traction, no normal flow) at the core–mantle and top surface boundaries.

Successful solution of realistic mantle flow problems must overcome a number of computational challenges due to the severe nonlinearity, heterogeneity, and anisotropy of earth’s rheology. Nonlinear behavior at narrow plate boundary regions influences the motion

of whole plates at continental scales, resulting in a wide range of spatial scales. Crucial features are highly localized with respect to earth’s radius (~ 6371 km), including plate thickness of order ~ 50 km and shear zones at plate boundaries of order ~ 5 km. Desired resolution at plate boundaries is below ~ 1 km. However, a mesh of earth’s mantle with uniform resolution of 0.5 km would result in $\mathcal{O}(10^{13})$ degrees of freedom (DOF), which would be prohibitive for models with such complexity. Thus adaptive methods are essential. Six orders of magnitude viscosity contrast is characteristic of the shear zones at plate boundaries, yielding sharp viscosity gradients and leading to severe ill-conditioning. Furthermore, the viscosity’s dependence on a power of the second invariant of the strain rate tensor and plastic yielding phenomena lead to severely nonlinear behavior.

Overcoming major obstacles in adaptivity [2], we advanced toward these challenges with global models having 20 km thick plate boundaries, nonlinear viscosity, and yielding, and in turn demonstrated an unanticipated level of coupling between plate motion and the deep mantle [3], bounds on energy dissipation within plates [4], and rapid motion of small tectonic plates adjacent to large ones [5]. Nevertheless, such models did not close the gap between the fine-scale (~ 1 to 10 km) patterns of earthquakes, stress, and topography along plate boundaries with plate motions. Narrowing the local-to-global divide is essential for extracting the key observations allowing one to reach a new understanding of the physics of earth processes (Figure 1).

The central computational challenge to doing so is to design implicit solvers and implementations for high-resolution realistic mantle flow models that can handle the resulting extreme degrees of nonlinearity and ill-conditioning, the wide ranges of length scales and material properties, and the highly adapted meshes and required advanced discretizations, while also scaling to the $\mathcal{O}(10^6)$ cores characteristic of leadership class supercomputers. While the conventional view is that this goal is impossible, we demonstrate that with a careful redesign of discretization, algorithms, solvers, and implementation, this goal is indeed possible. These advances open the door to merging two fundamental geophysical questions: the origin of great earthquakes and the balance of forces driving plate motions.

II. STATE OF THE ART

Earth’s mantle convection is one of a large number of complex PDE problems that require implicit solution

on extreme-scale systems. The complexity arises from the presence of a wide range of length scales and strong heterogeneities, as well as localizations and anisotropies. Complex PDE problems often require aggressive adaptive mesh refinement, such as that provided by the parallel forest-of-octree library p4est [6] that originated in our group. They also often require advanced discretizations, such as the high-order, hanging-node, mixed continuous-velocity/discontinuous-pressure element pair employed here. The physics complexities combined with the discretization complexities conspire to present enormous challenges for the design of solvers that are not only algorithmically optimal, but also scale well in parallel. These challenges are well documented in a number of blue ribbon panel reports (e.g., [7]).

In our context of time-independent non-Newtonian flows, *implicit solvers* means a combination of nonlinear and linear solvers and preconditioners. We employ Newton’s method, the gold standard for nonlinear solvers. It can deliver asymptotic quadratic convergence, independent of problem size, for many problems. However, differentiating complex constitutive laws such as (1c) to obtain the linearized Newton operator creates an even more complex system to be solved. Combining the Newton method with an appropriately truncated Krylov linear solver permits avoidance of *oversolving* far from the region of fast convergence [8]. The crucial point is then the preconditioner, which must simultaneously globalize information to maximize algorithmic efficiency, while localizing it to maximize parallel performance. For preconditioning, we target multilevel solvers, which are algorithmically optimal for many problems (i.e., they require $\mathcal{O}(n)$ work, where n is the number of unknowns) and parallelize well (requiring $\mathcal{O}(\log n)$ depth), at least for simple elliptic PDE operators.

The state of the art in extreme-scale multilevel solvers is exemplified by the Hybrid Hierarchical Grids (HHG) geometric multigrid (GMG) method [9], the GMG solver underlying the UG package [10], the algebraic multigrid (AMG) solver BoomerAMG from the hypre library [11], the multilevel balancing domain decomposition solver in FEMPAR [12], and the AMG solver for heterogeneous coefficients from the DUNE project [13]. These multigrid solvers have all been demonstrated to scale up to several hundred thousand cores (458K cores in some cases), but only for constant coefficient linear operators, uniformly-refined meshes, and low-order discretizations (with the exception of the DUNE solver, which has demonstrated scalability on a problem with heterogeneous coeffi-




Figure 2: Comparison of algorithmic performance of conventional state-of-the-art (*dashed lines*) vs. new (*solid lines*) Stokes solver for a sequence of increasingly difficult problems (indicated by colors), reflecting increasingly narrower plate boundary regions.

cients but otherwise with uniform and low order grids). The complex PDE problems we target—characterized by advanced high-order discretizations, highly-locally adapted meshes, extreme (six orders-of-magnitude variation) heterogeneities, anisotropies, and severely nonlinear rheology—are significantly more difficult. We are not aware of any solver today that is capable of solving such problems at large scale, with algorithmic and parallel efficiency.

When $\mathcal{O}(10^5)$ cores and beyond are needed for implicit solution of such complex PDE problems, the usual approach has been to retreat to algorithmically suboptimal but easily-parallelizable solvers (such as explicit or simply-preconditioned implicit). This is clearly not a tenable situation, and the performance gap between optimal and suboptimal solvers only increases as problems grow larger. Thus our goal here is to present an implicit solver (significantly going beyond our previous work [14], [15], [2]) that delivers optimal algorithmic complexity while scaling with high parallel efficiency to the full size of leadership-class supercomputers for the class of complex PDE problems targeted here, with particular application to our driving global mantle convection problem.

Figure 2 illustrates the power of algorithmically optimal solvers for our mantle convection problem. The curves show the reduction in residual as a function of Krylov iterations for a sequence of increasingly difficult problems (different colors). The dashed curves represent a contemporary, well-regarded solver, such as that found in the state-of-the-art community mantle convection code ASPECT [16]. This combines AMG to precondition the (1,1) block of the Stokes system along with a diagonal mass matrix approximation of the (2,2) Schur complement. Our new solver (see next section) combines a sophisticated hybrid spectral-geometric-algebraic multigrid (HMG) along with a novel HMG-preconditioned improved Schur complement approximation. The mas-

sive enhancement in algorithmic performance (over 4 orders of magnitude lower residual for the same number of iterations) seen in the figure is due to the improvement of the Schur complement. This is what makes the solution of the high-fidelity mantle flow models we are targeting tractable. It increases however the algorithmic complexity, but as we will see, we are still able to obtain excellent scalability out to 1.5M cores, to go with the several orders of magnitude improvement in run time. Key to achieving this scalability is: (i) avoiding AMG setup/communication costs with a spectral and geometric multigrid approach; (ii) eliminating AMG’s requirement for matrix assembly and storage for differential operators and intergrid transfer operations.

III. INNOVATIVE CONTRIBUTIONS

A. Summary of contributions

Hybrid spectral-geometric-algebraic multigrid (HMG). *High-order* discretizations on *locally refined* meshes for implicit problems with *extreme variations* in coefficients pose challenge for extreme-scale PDE solvers. We develop a multigrid scheme based on *matrix-free* operators that does not require collective communication and repartitions meshes at coarse multigrid levels. The latter is achieved using a hierarchy of MPI communicators for point-to-point communication. In this way, we obtain optimal time-to-solution.

Preconditioner. The Schur complement approximation in our solver is known to be critical for problems with extreme variations in the coefficients. We propose a new *HMG-based* approach for preconditioning the Schur complement of the nonlinear Stokes equations. It extends a Schur complement method based on discrete arguments which limited its application to AMG. Since AMG is difficult to scale to millions of cores and has a significant memory footprint, we have developed an HMG method based on a PDE operator that mimics the algebraic operator occurring in the preconditioner.

Nonlinear solver. For the first time, a *grid-continuation, inexact Newton-Krylov method* is used for a realistic and severely nonlinear rheology over the entire earth. The nonlinearity originates from power law shear thinning, viscosity bounds, and plastic yielding [17]. This new method enables us to simulate the global instantaneous mantle flow in the entire earth, with *unprecedented accuracy*.

B. Algorithm overview

We employ an inexact Newton-Krylov method for the nonlinear Stokes equations (1), i.e., we use a sequence of linearizations of (1) and approximately solve the resulting linearized systems using a preconditioned Krylov method. The design of the preconditioner required the majority of the algorithmic innovations, but the nonlinear solver components needed careful consideration as well. In particular, we define the rheology (1c) such that it incorporates bounds for the viscosity in a differentiable manner, permitting the use of Newton’s method. To compute a Newton update $(\tilde{\mathbf{u}}, \tilde{p})$, we find the (inexact) solution of the linearized Stokes system,

$$\begin{aligned} -\nabla \cdot \left[\mu' (\nabla \tilde{\mathbf{u}} + \nabla \tilde{\mathbf{u}}^\top) \right] + \nabla \tilde{p} &= -\mathbf{r}_{\text{mom}}, \\ \nabla \cdot \tilde{\mathbf{u}} &= -r_{\text{mass}}, \end{aligned} \quad (2)$$

with

$$\mu' = \mu \mathbf{I} + \dot{\varepsilon}_{II} \frac{\partial \mu}{\partial \dot{\varepsilon}_{II}} \frac{(\nabla \mathbf{u} + \nabla \mathbf{u}^\top) \otimes (\nabla \mathbf{u} + \nabla \mathbf{u}^\top)}{\|(\nabla \mathbf{u} + \nabla \mathbf{u}^\top)\|_F^2}, \quad (3)$$

where the current velocity and pressure are \mathbf{u} and p , respectively, and the residuals of the momentum and mass equations appear on the right-hand side of (2). Note that what plays the role of viscosity in the Newton step is an anisotropic fourth-order tensor (3).

We discretize earth’s mantle using locally adaptively refined hexahedral meshes. Extreme local refinement is critical to resolve plate boundaries down to a few hundred meters, while away from these regions significantly coarser meshes can be used that still capture global-scale behavior. Parallel adaptive forest-of-octrees algorithms, implemented in the p4est parallel AMR library, are used for efficient parallel mesh refinement/coarsening, mesh balancing, and repartitioning [2], [6], [18]. In (2), the velocity is discretized with high-order, non-conforming, continuous nodal finite elements of polynomial order $k \geq 2$, and the pressure with discontinuous modal elements of order $k - 1$. This velocity-pressure pairing yields optimal asymptotic convergence with decreasing mesh element size and conserves mass locally at the element level. It is provably inf-sup stable and thus avoids stabilization terms, which can degrade the accuracy of the solution, especially in mantle convection simulations.

This discretization of the Newton step results in an extremely ill-conditioned algebraic system with up to hundreds of billions of unknowns, which requires a preconditioned Krylov iterative method. Such a Krylov method needs only the application of the left hand side

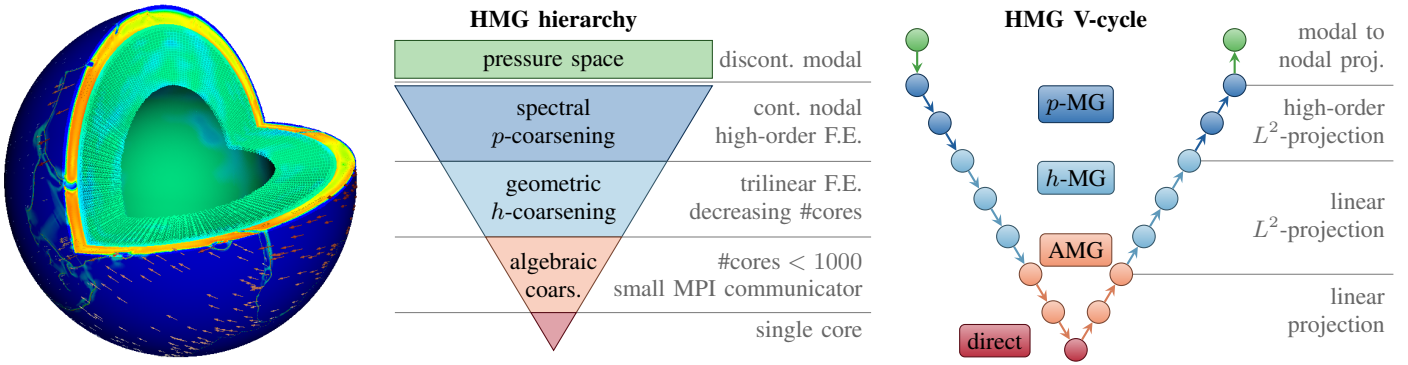


Figure 3: *Left image:* visualization of a simulation and part of the computational domain showing the adaptively refined mesh. The *color coding* illustrates the effective mantle viscosity and the *arrows* depict the motion of the tectonic plates. *Central diagram:* illustration of multigrid hierarchy. From top to bottom, first, the multigrid levels are obtained by spectral coarsening. Next, the mesh is geometrically coarsened and repartitioned on successively fewer cores to minimize communication. Finally, AMG further reduces problem size and core count. The multigrid hierarchy used in the Schur complement additionally involves smoothing in the discontinuous modal pressure space (*green*). *Right diagram:* the multigrid V-cycle consists of smoothing at each level of the hierarchy (*circles*) and intergrid transfer operators (*arrows downward* for restriction and *arrows upward* for interpolation). To enhance efficacy of the the V-cycle as a preconditioner, different types of projection operators are employed for these operators depending on the phase within the V-cycle.

operator in (2) to vectors, which we implement in a matrix-free fashion using elemental loops. We exploit the tensor-product structure of the element-level basis functions, resulting in a reduced number of operations [19]. We use GMRES as the Krylov solver, with right preconditioning based on the upper triangular block matrix:

$$\underbrace{\begin{bmatrix} \mathbf{A} & \mathbf{B}^\top \\ \mathbf{B} & \mathbf{0} \end{bmatrix}}_{\text{Stokes operator}} \underbrace{\begin{bmatrix} \tilde{\mathbf{A}} & \mathbf{B}^\top \\ \mathbf{0} & \tilde{\mathbf{S}} \end{bmatrix}^{-1}}_{\text{preconditioner}} \begin{bmatrix} \tilde{\mathbf{u}} \\ \tilde{\mathbf{p}} \end{bmatrix} = \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{bmatrix}. \quad (4)$$

Here, approximations of the inverse of the viscous block, $\tilde{\mathbf{A}}^{-1} \approx \mathbf{A}^{-1}$, and the inverse of the Schur complement, $\tilde{\mathbf{S}}^{-1} \approx (\mathbf{B}\mathbf{A}^{-1}\mathbf{B}^\top)^{-1}$, are required, where \mathbf{B} and \mathbf{B}^\top denote the discrete divergence and gradient matrices. This particular combination of Krylov method and preconditioner type is known to converge in only two iterations for optimal choices of $\tilde{\mathbf{A}}^{-1}$ and $\tilde{\mathbf{S}}^{-1}$ [20].

The inverse of the viscous block $\tilde{\mathbf{A}}^{-1}$ is approximated by a multigrid V-cycle, as detailed below. For the inverse of the Schur complement $\tilde{\mathbf{S}}^{-1}$, we use an improved version of the Least Squares Commutator method [21], [22]: $\tilde{\mathbf{S}}^{-1} = (\mathbf{B}\mathbf{D}^{-1}\mathbf{B}^\top)^{-1}(\mathbf{B}\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1}\mathbf{B}^\top)(\mathbf{B}\mathbf{D}^{-1}\mathbf{B}^\top)^{-1}$. It has been demonstrated to be robust with respect to extreme viscosity variations, as shown in Table I. This approach requires approximating the inverse of $\mathbf{B}\mathbf{D}^{-1}\mathbf{B}^\top$, where $\mathbf{D} := \text{diag}(\mathbf{A})$.

Multigrid is a natural choice to invert this matrix. However, AMG would require matrix assembly and an expensive setup. The problem lies in the particular coupling between matrices \mathbf{B} and \mathbf{B}^\top . Computing the

product matrix results in large communication requirements and in a large number of nonzero entries. The number of nonzero entries of the product matrix increase similarly as when squaring the matrix of a discrete Laplacian. We thus developed an approach that uses the analogy between $\mathbf{B}\mathbf{D}^{-1}\mathbf{B}^\top$ and an anisotropic, variable-coefficient elliptic PDE operator. This operator is discretized with continuous, k th order, nodal finite elements. This continuous, nodal Poisson operator, which we call \mathbf{K} , is then inverted with an HMG V-cycle plus additional smoothing steps in the discontinuous modal pressure space. For smoothing in the pressure space we compute and store only the diagonal entries of $\mathbf{B}\mathbf{D}^{-1}\mathbf{B}^\top$, which requires no communication.

As illustrated in Figure 3, our HMG method is divided in four stages (or even five stages in the pressure Poisson case). This hybrid multigrid setup sits at the core of our nonlinear solver and thus a careful design of the intergrid transfer operators was critical for efficiency and performance. Our hybrid multigrid method combines high-order L^2 -restrictions/interpolations, uses the full fourth-order tensor coefficient in the Newton step (3) on all levels, and employs Chebyshev-accelerated point-Jacobi smoothers. This results in optimal algorithmic multigrid performance, i.e., iteration numbers are independent of mesh size and discretization order, and are robust with respect to the highly heterogeneous coefficients (six orders of magnitude viscosity and nine orders of magnitude Poisson coefficient contrast) occurring in the simulation of mantle flow with plates (see Section IV).

C. Implementation and optimization

From a high-level perspective, the challenge of a parallel multigrid implementation is to balance the performance of two critical components: (i) application of differential operators during smoothing, commonly referred to as MatVecs, and (ii) intergrid transfer operators that perform restriction and interpolation between multigrid levels. Further, this balance has to be maintained as the number of cores grows to extreme scales. Both MatVecs and intergrid operators rely on point-to-point communication such that optimizing the runtime of one deteriorates the performance of the other. In the case of our complex mantle flow solver, we deal with four different kinds of MatVecs (viscous stress \mathbf{A} , divergence/gradients $\mathbf{B}/\mathbf{B}^\top$, continuous nodal Poisson operator \mathbf{K} , and Stokes operator) and six different intergrid operators (restriction and interpolation for each of: modal to nodal projection, p -projection in spectral multigrid, and h -projection in geometric multigrid). Optimization efforts have to target all of these operators to be successful. Additionally, this task is highly non-trivial, since the HMG V-cycle has to be performed on unstructured, highly locally-adapted meshes.

In order to obtain optimal load balance for MatVecs during the V-cycle, we repartition the coarser multigrid levels uniformly across the cores and gradually reduce the size of the MPI communicators as we progress through the coarser levels. The reduction of the MPI size is done such that neither MatVecs nor intergrid transfer operations become a bottleneck at large scale. Moreover, point-to-point communication is overlapped with computations for optimal scalability. No collective communication is used in the V-cycle. The HMG setup cost is minimized with a matrix-free approach for differential and intergrid operators, which additionally produces a lightweight memory footprint.

These key principles were at the foundation of our extreme-scale multigrid implementation. Further improvements of time-to-solution and performance were carried out in a number of successive optimization steps (see Figure 4a). Overall, we decreased the time to solution for the targeted hardware architecture (see Section V) by a factor of over 1000 and increased performance on a compute node by a factor of ~ 200 . With this performance, our complex mantle flow solver as a whole, including spectral, geometric and algebraic multigrid phases on highly adaptively refined meshes, is as performant as a routine for sparse matrix-vector




Figure 4: (a) Performance improvement and time-to-solution reduction over a sequence of optimization steps (time is normalized by GMRES iterations per 1024 BG/Q nodes per billion DOF). *Pt. A* is base performance before optimization. *Pt. B:* reduction of blocking MPI communication. *Pt. C:* minimizing integer operations in inner MatVec for-loops and reducing the number of cache misses. *Pt. D:* computation of derivatives by applying precomputed CSR-matrices at the element level and SIMD vectorization. *Pt. E:* OpenMP threading of major loops in MatVecs. *Pt. F:* MPI communication reduction, overlapping with computations, and OpenMP threading in intergrid operators. *Pt. G:* low-level optimization of finite element kernels via improving flop-byte ratio and consecutive memory access, and better pipelining of floating point operations. *Pt. H:* various low-level optimizations including enforcement of boundary conditions and interpolation of hanging finite element nodes. (b) BG/Q node roofline model (theoretical peak performance) and SpMV performance with max flop-byte ratio of 0.25 [24].

multiplications. This is supported by the roofline model analysis [23], from which we obtain an optimal performance of approximately 8 GFlops for sparse MatVecs (Figure 4b). Note that implicit solvers for PDEs inherently exhibit a sparsity structure and hence performance will always be memory-bound. This argument shows that the performance of our memory-bound solver is close to what is optimally achievable. This is further supported by numerical results in Section VI.

IV. EXPERIMENTAL SETUP AND VERIFICATION OF OPTIMALITY AND ROBUSTNESS OF THE SOLVER

In this section we describe the physical problem and solver parameters used to carry out the parallel performance analysis in Section VI. We also present results of tests of optimal algorithmic scalability and robustness of the solver which together with parallel scalability demonstrate overall scalability of the solver.

The important physical parameter that determines the difficulty of the problem is the viscosity field. In our subsequent performance analysis, we use real earth data to generate a physically realistic representation of viscosity (which is a function of T , w , and $\dot{\epsilon}_{II}$). The viscosity varies over six orders of magnitude globally. However, what makes realistic mantle flow problems even more highly ill-conditioned and nonlinear is the

extremely thin layer in which this contrast develops. The viscosity drops by six orders of magnitude within a thin layer between two plates (the *plate boundary*).

To assess solver robustness and algorithmic scalability, we generate plate boundaries down to a thickness of 5 km and a factor of 10^6 viscosity drop over 7 km. For the weak and strong scalability measurements, the 10^6 factor viscosity drop occurs within just 3 km. Since tectonic plates (the largest surface structures) are 2,000–14,000 km across, and earth’s circumference is 40,075 km, this results in a very wide range of length scales of interest. To capture the viscosity variation, the mesh is refined to ~ 75 m local resolution in our largest simulations, resulting in a mesh with 9 levels of refinement. For all performance results, we use a velocity discretization with polynomial order $k = 2$.

The cost of solving a nonlinear earth mantle flow problem is dominated by the cost of the linear solve in each Newton step (2). The cost of a linear solve is determined by the number of MatVecs and HMG inter-grid operations. MatVecs are encountered in the Krylov method and in the HMG smoothers. In all subsequently reported performance results, we use three smoothing iterations for pre- and post-smoothing within the HMG V-cycle for both the viscous block $\tilde{\mathbf{A}}^{-1}$ and in the Schur complement $\tilde{\mathbf{S}}^{-1}$, which amounts to three V-cycles per application of the Stokes preconditioner. Therefore each GMRES iteration has the same cost and it is sufficient to compare the number of GMRES iterations.

A. Robustness of Stokes solver

The robustness of the HMG preconditioner for the Stokes system (1) is assessed by observing the number of GMRES iterations (see Table I) required for convergence, while decreasing the thickness of the plate boundaries. This increases the range of length scales in the problem and the resulting nonlinearity and ill-conditioning. Our mesh refinement algorithm, which is based on the norm of the viscosity gradients as well as the magnitude of ε_{II} , locally refines the mesh to resolve the extreme viscosity variations. This results in an overall increase in the number of DOF. The third and fourth columns in Table I demonstrate the robustness of the solvers for the (1,1) Stokes block and for the complete Stokes solver. The GMRES iterations are seen to scale independently of the plate boundary thickness and thus viscosity gradient. Similar independence of viscosity is observed for nonlinear iterations.

Table I: Robustness with respect to plate boundary thickness of HMG-preconditioned GMRES solver for the (1,1) block of Stokes and the linear (full) Stokes solver. Number of GMRES iterations to reduce the residual by a factor of 10^{-6} is reported.




Plate boundary thickness [km]	DOF [$\times 10^9$]	GMRES iterations to solve $\mathbf{A}\mathbf{u} = \mathbf{f}$	GMRES iterations to solve Stokes
15	1.16	115	461
10	1.41	129	488
5	3.01	123	445

B. Algorithmic scalability

Algorithmic scalability, i.e., the independence of the solver iterations from the resolution of the mesh, is critical for overall scalability of implicit solvers. To study algorithmic scalability, we consider a nonlinear mantle flow problem with one subducting slab. The plate boundary region between the subducting plate and the overriding plate has a thickness of 5 km. We refine the mesh locally in the regions of highest viscosity variations by tightening the refinement criteria. Thus the total number of DOF grows slowly, though significantly greater resolution is obtained in these regions. The required numbers of linear and nonlinear iterations are shown in Table II, where the cost of the nonlinear solver is measured by the total number of GMRES iterations across nonlinear iterations. As can be seen, the linear solver requires a number of iterations that is largely independent of the resolution of the problem.

We have demonstrated how the combination of our preconditioner and linear and nonlinear solver yields an implicit method whose number of iterations scales independent of model fidelity. Here, fidelity is understood as the resolution of the mesh with finite element discretization and the size of the smallest-scale features, which are the plate boundary regions. This results in an algorithmically optimal method, despite the severely nonlinear rheology, high viscosity gradients, effective anisotropy, and large heterogeneities. Moreover, the cost of the solver is reduced by adaptive mesh refinement, which reduces the number of DOF—in this

Table II: Optimal algorithmic scalability of inexact Newton-Krylov method for solving a nonlinear mantle flow problem with one subducting slab and 5 km plate boundary. Simulation cost expressed in *total* number of GMRES iterations is largely independent of the maximal resolution of the adaptive mesh (10^{-7} Newton residual reduction used as stopping criterion). A two times higher resolution increases the DOF of the adaptively refined mesh only by about a factor of 2–3. In contrast, the factor would be eight with uniform refinement.

Max level of refinement	Finest resolution [m]	DOF [$\times 10^6$]	Newton iterations	GMRES iterations
10	2443	0.96	14	1408
11	1222	2.67	18	1160
12	611	5.58	21	1185
13	305	11.82	21	1368
14	153	36.35	27	1527

case by four orders of magnitude, from the $\mathcal{O}(10^{13})$ needed for a uniform mesh of earth’s mantle, to just the $\mathcal{O}(10^9)$ required here using aggressive refinement. Further reducing the number of DOF are the third-order accurate finite elements employed here, along with a mass-conserving discretization. The algorithmic scalability and the greatly-reduced number of DOF exhibited by our solver are critical for the overall goal of reducing time-to-solution (for a given accuracy). The remaining component is parallel scalability, which we study next.

V. SYSTEMS AND MEASUREMENT METHODOLOGY

The target architecture in this work is the IBM BlueGene/Q¹ (BG/Q) supercomputer [25]. Table III summarizes size and peak performance of several systems we used. The smaller systems were used for testing, optimization, scaling, and full science runs. The largest runs have been performed on the Sequoia supercomputer at the Lawrence Livermore National Laboratory (LLNL). Sequoia consists of 96 IBM Blue Gene/Q racks, reaching a theoretical peak performance of 20.1 PFlops/s. Each rack consists of 1024 compute nodes, each hosting an 18 core A2 chip that runs at 1.6 GHz. Of these 18 cores, 16 are devoted to computation, one for the lightweight O/S kernel, and one for redundancy. Every core supports 4 H/W threads, thus, in total Sequoia has 1,572,864 cores and can support up to 6,291,456 H/W threads. The total available system memory is 1.458 PBytes. BG/Q nodes are connected by a five-dimensional (5-D) bidirectional network, with a network bandwidth of 2 GBytes/s for sending and receiving data. Each BG/Q rack features

¹IBM and Blue Gene/Q are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies.

Table III: Blue Gene/Q supercomputers

	Racks	Cores	H/W threads	Peak [PFlops/s]
AMOS	5	81,920	327,680	1.0
Vulcan	24	393,216	1,572,864	5.0
JUQUEEN	28	458,752	1,835,008	5.8
Sequoia	96	1,572,864	6,291,456	20.1

dedicated I/O nodes with 4 GBytes/s I/O bandwidth. The system implements optimized collective communication and allows specialized tuning of point-to-point communication. We obtained all timing and performance measurements by means of the IBM HPC Toolkit for BG/Q. The toolkit retrieves performance information about the processor, memory hierarchy, and interconnect. Power measurements are available through hardware sensors for each node board (accommodating 32 nodes) at a rate of 2 samples per second [26]. All runs involved double-precision arithmetic and the code is compiled using the IBM XL C compilers for BG/Q, version 12.1.10.

VI. PERFORMANCE RESULTS

A. Weak and strong scalability

We present weak and strong scalability results on the Vulcan and Sequoia BG/Q supercomputers from 1 rack with 16,384 cores up to 96 racks with 1,572,864 cores. Scalability measurements corresponding to 1, 2, and 4 racks were obtained on Vulcan, whereas the remaining runs on 8–96 racks were performed on Sequoia.

The cost of our large-scale nonlinear mantle convection simulations is overwhelmingly dominated by the cost of the GMRES iterations during a linear Stokes solve (everything else, including setup and I/O, is negligible). These GMRES iterations include HMG V-cycles for the (1,1) Stokes block and in the Schur complement approximation, as explained earlier. For the extreme-scale runs on Sequoia, we had limited access to the system, which allowed just 10 representative GMRES iterations. However, we illustrate the influence of I/O and setup costs by extrapolating the number of GMRES iterations to those expected for a full nonlinear solution. Note that we did observe that the setup time for HMG is largely bounded independent of the number of cores.

The main result is the weak scalability shown in Figure 5. The solver maintained 97% parallel efficiency (red curve) over a 96-fold increase in problem size, from 16K to 1.5M cores of the full Sequoia system. The largest problem involved 602 billion DOF. If we take into




Figure 5: Weak scalability results on Vulcan and Sequoia from 1 to 96 racks. Performance is normalized by time and number of GMRES iterations. Numbers along the graph lines indicate efficiency w.r.t. ideal speedup (efficiency baseline is the 1 rack result). We report both the weak scalability for the linear solver only (red) and for projected total runtime of a nonlinear solve (green). The largest problem size on 96 racks has 602 billion DOF.

account the setup time of the problem, including I/O of input data and solver setup time, we would still arrive at a weak scalability efficiency of 96% (green curve) for total runtime of a nonlinear simulation, demonstrating the negligibility of I/O and setup time. The I/O for writing output data has to be performed only once at the end of a nonlinear solve. The problem sizes used in the weak scalability runs would produce ~ 8.5 GBytes of output per BG/Q I/O node. With an I/O bandwidth of 4 GBytes/s we can also consider the writing of the output to be negligible for overall runtime (note that we did not output solution fields, since the full nonlinear simulation could not be run to completion due to limited access). The negligible time for I/O and problem setup stem from the advantages of adaptive implicit solvers: adaptivity results in the problem itself being generated online as part of the solver; implicit means that fewer outputs/checkpoints would be required.

In Figure 6, we show strong scalability results for a mantle convection simulation with 8.3 billion DOF. Starting from one rack with 16,384 cores (granularity of 506K DOF/core), we achieve a 32-fold speedup on 96 racks with 1,572,864 cores (granularity of 5K DOF/core), indicating 33% solver efficiency in strong scalability, an impressive number considering the coarse granularity of the largest problem.

Contrary to conventional wisdom, this shows that algorithmically optimal implicit finite element solvers for severely nonlinear, ill-conditioned, heterogeneous, indefinite PDEs can be designed to scale to $\mathcal{O}(10^6)$ cores.

B. Node performance analysis

The performance results on BG/Q compute nodes further support our scalability results. The top pie charts of




Figure 6: Strong scalability results on Vulcan and Sequoia from 1 to 96 racks. Numbers along the graph lines indicate efficiency with respect to ideal speedup (efficiency baseline is the 1 rack result). We report both the strong scalability for the linear solver only (red) and for projected total runtime of a nonlinear solve (green).

Figure 7 decompose the overall runtime into the largest contributors. We can observe that the (highly optimized) matrix-free apply routines dominate with 80.6% in the 1 rack case. Furthermore, their portion remains very stable with 78% on 96 racks. This result demonstrates a key component of a highly scalable, parallel multigrid implementation. The percent runtime for intergrid transfer operations is low compared to MatVecs and stays low even at 1.5 million cores. Hence, we have achieved a balance between MatVecs and intergrid operations that results in nearly optimal scalability.

MatVecs represent the portion of the code where the maximal performance in terms of flops can be achieved. With their dominance in runtime we are able to increase total performance close to its maximum. That way our implementation is performing at the limits of the roofline model as predicted in Figure 4b, and this is achieved even at extreme scales of $\mathcal{O}(10^6)$ cores.

C. MPI communication analysis

Figure 8 summarizes MPI communication time measured during weak and strong scalability runs: tasks with minimum, median, and maximum communication time are displayed. Indeed, for weak scalability, we clearly observe that percentage of time spent in MPI communication remains nearly constant relative to runtime (Figure 8a). This contributes to the nearly perfect scalability results presented in Figure 5. The increase in median and maximum communication time in the 64 racks case can be justified by the lack of 5-D torus connectivity in that particular configuration (due to specific job partitioning). Another reason can be found in a more aggressive repartitioning of coarser multigrid levels, which leaves a greater amount of cores idle during a short period of time in the V-cycle. This is suggested by the higher percentage of MPI_Waitall time on 64 racks in

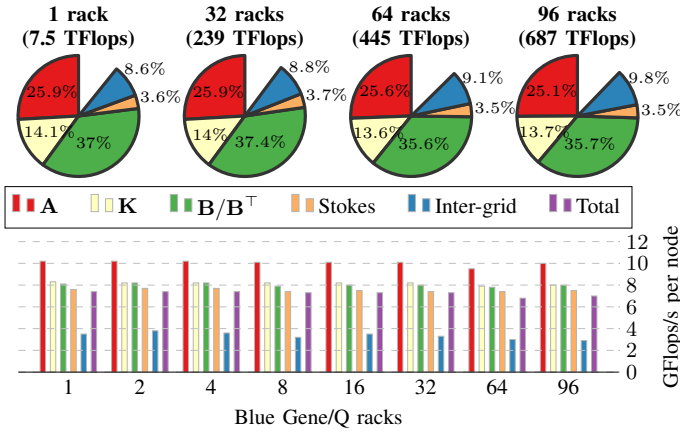


Figure 7: Analysis of MatVecs and intergrid operators within the Stokes solves of the weak scalability runs on Vulcan and Sequoia. Pie charts show fraction of time in each routine, while the histograms show corresponding average GFlops/s per BG/Q node. The symbols denote MatVecs for viscous stress \mathbf{A} , continuous, nodal Poisson operator \mathbf{K} , and divergence/gradient \mathbf{B}/\mathbf{B}^T (see Section III). The empty slices in the pie charts consist of all other routines with generally low GFlops/s per node (e.g., GMRES orthogonalization, null space projections).

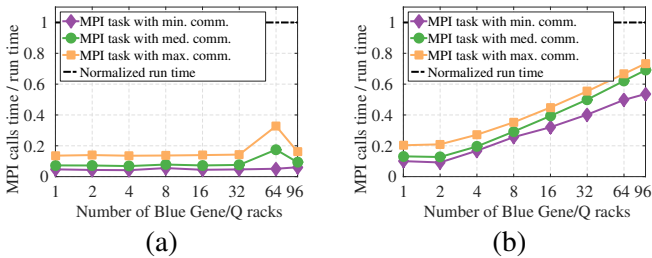


Figure 8: MPI communication time relative to total runtime for (a) weak scalability and (b) strong scalability on Vulcan and Sequoia.

Figure 9. However, this does not need to affect scalability in a negative way since fewer cores may perform the same task quicker because of higher granularity of DOF.

For the strong scalability runs, we observe a gradual increase of relative MPI communication time (Figure 8b), as is expected for implicit solvers. Note that the increase begins only at 4 racks. Communication time exceeds 50% of overall runtime only at about 1 million cores. At its maximum, communication time is still below 30%.

D. Energy consumption analysis

Finally, we analyze the energy efficiency of the scalability runs. As expected, during weak scalability energy consumption increases linearly with the amount of used resources (Figure 10a). With an estimated cost of \$0.06 per kWh [27], the energy cost per GMRES iteration on

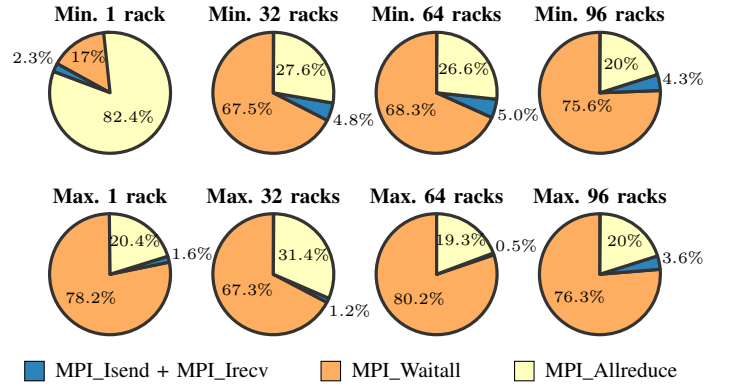


Figure 9: MPI routine communication time for the weak scalability runs on Vulcan and Sequoia supercomputers.

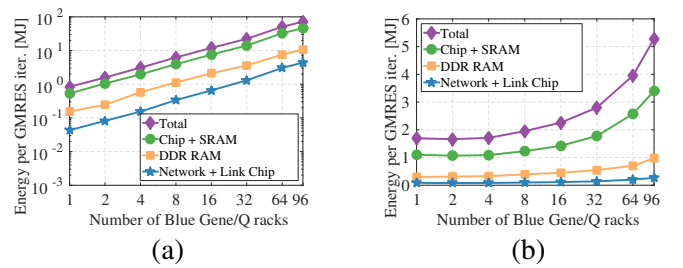


Figure 10: Scalability of energy (excluding cooling) on Vulcan and Sequoia: (a) weak scalability, (b) strong scalability.

96 racks is nearly \$1.20 (excluding cooling). On the other hand, the loss of strong scalability on the full size of the machine is reflected in the energy usage (Figure 10b). The 33% speedup efficiency on 96 racks (see Figure 6) is combined with a 32% energy efficiency; in other words, power consumption per node does not change significantly, and energy efficiency is mainly driven by time-to-solution.

VII. IMPLICATIONS FOR MANTLE FLOW MODELING

Building on algorithmic innovations for implicit solvers for complex PDEs described in this paper, we are able to address the depth and distribution of oceanic trenches—the most extreme topographic features on earth’s surface—for the first time. Trench depth reflects both the downward pull from plate-driving forces [28] and the variable resistance associated with seismic coupling from great earthquakes [29]. We forward-predict the width (~ 50 km) and depth (~ 10 km) of oceanic trenches on a global scale while predicting plate motions (Figure 11). The simultaneous prediction of these quantities—large-scale flow and fine-scale stress at plate boundaries—in a model with realistic, nonlinear rheology employing scalable, robust solvers opens new

directions for geophysical research. Solver robustness to plate boundary thickness is crucial, as can be seen in Figure 12, where we observe a great sensitivity of the simulation outcome (in terms of plate velocities) to the thickness of plate boundary regions. The scalable solver presented in this paper, in combination with adjoints, which are a byproduct of the Newton solver, will allow systematic inference of uncertain parameters in global mantle flow systems with tectonic plates. For regional mantle models that are functionally equivalent to the global computation presented here, we have recently illustrated a systematic inference approach for the nonlinear constitutive parameters n and τ_{yield} , and plate coupling factors $w(x)$, for several subduction zones [30]. Adjoint-based inversions will require thousands of forward model solutions, so that availability of a scalable implicit solver such as that described here is paramount.

Bringing observations on topography (trench depth), plate motions, and others into a global inversion will allow the merging of two distinct geophysical approaches at different scales addressing different questions. First, what is the degree of coupling associated with great earthquakes? In particular, we seek to determine whether that coupling is due to the frictional properties of the incoming plate or the magnitude of normal stress across the fault driven by tectonic processes [31], [32]. The second question concerns the forces driving and resisting global plate motion and the degree to which inter-plate coupling governs plate motions [4], [33], [34]. These questions have eluded solution over the past three decades because of their intimate coupling. By allowing us to bridge the local-to-global scales, with modern data sets, will arguably allow us to make an important leap toward the simultaneous solution of two of the most fundamental questions in earth sciences.

ACKNOWLEDGMENTS

We wish to acknowledge the contributions of W. Scott Futral (LLNL) and Roy Musselman (IBM), who were instrumental in helping us achieve the scaling results on Sequoia. Their contributions came after SC’s July deadline for finalizing the author list had passed, but they should be regarded as co-authors.

We wish to offer our deepest thanks to Lawrence Livermore National Laboratory, Jülich Supercomputing Center, RPI Center for Computational Innovation, and Texas Advanced Computing Center for granting us the computing resources required to prove the pioneering nature of this work. This research was partially




Figure 11: Results of nonlinear mantle flow simulation at the earth surface: View centered on 180°W (left) and 90°W (right) showing north–westward motion of the Pacific Plate (black arrows) and the total normal stress field (color coded). This stress is proportional to the dynamic topography and for the first time we are able to forward predict narrow (~ 50 km in width) ocean trenches (narrow lines with dark blue color) along plate boundaries in a global model with plate motions.

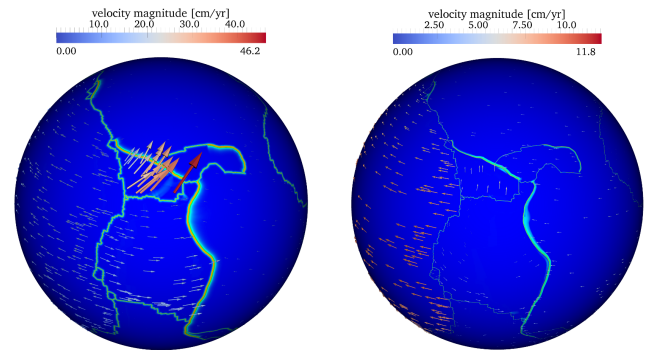


Figure 12: Comparison of earth plate velocities of a low-fidelity model (left) and a high-fidelity model (right) with thinner plate boundaries. Significant sensitivity of velocities of the Cocos Plate (in center) are observed. This illustrates the importance of the solver’s ability to handle a wide range of values for plate boundary thickness.

supported by NSF grants CMMI-1028889 and ARC-0941678 and DOE grants DE-FC02-13ER26128 and DE-FG02-09ER25914 as well as the EU FP7 EXA2GREEN and NANOSTREAMS projects. We also thank Carsten Burstedde for his dedicated work on the p4est library.

REFERENCES

- [1] D. Depaolo, T. Cerling, S. Hemming, A. Knoll, F. Richter, L. Royden, R. Rudnick, L. Stixrude, and J. Trefil, “Origin and Evolution of Earth: Research Questions for a Changing Planet,” National Academies Press, Committee on Grand Research Questions in the Solid Earth Sciences, National Research Council of the National Academies, 2008.
- [2] C. Burstedde, O. Ghattas, M. Gurnis, T. Isaac, G. Stadler, T. Warburton, and L. C. Wilcox, “Extreme-scale AMR,” in *Proceedings of SC10*. ACM/IEEE, 2010.
- [3] L. Alisic, M. Gurnis, G. Stadler, C. Burstedde, and O. Ghattas, “Multi-scale dynamics and rheology of mantle flow with plates,” *Journal of Geophysical Research*, vol. 117, p. B10402, 2012.

- [4] G. Stadler, M. Gurnis, C. Burstedde, L. C. Wilcox, L. Alisic, and O. Ghattas, "The dynamics of plate tectonics and mantle flow: From local to global scales," *Science*, vol. 329, no. 5995, pp. 1033–1038, 2010.
- [5] L. Alisic, M. Gurnis, G. Stadler, C. Burstedde, L. C. Wilcox, and O. Ghattas, "Slab stress and strain rate as constraints on global mantle flow," *Geophysical Research Letters*, vol. 37, p. L22308, 2010.
- [6] C. Burstedde, L. C. Wilcox, and O. Ghattas, "p4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees," *SIAM Journal on Scientific Computing*, vol. 33, no. 3, pp. 1103–1133, 2011.
- [7] J. Dongarra, J. Hittinger, J. Bell, L. Chacón, R. Falgout, M. Heroux, P. Hovland, E. Ng, C. Webster, and S. Wild, "Applied mathematics research for exascale computing," Report of the DOE/ASCR Exascale Mathematics Working Group, 2014.
- [8] S. C. Eisenstat and H. F. Walker, "Choosing the forcing terms in an inexact Newton method," *SIAM Journal on Scientific Computing*, vol. 17, pp. 16–32, 1996.
- [9] B. Gmeiner, U. Rüde, H. Stengel, C. Waluga, and B. Wohlmuth, "Performance and scalability of Hierarchical Hybrid Multigrid solvers for Stokes systems," *SIAM Journal on Scientific Computing*, vol. 37, no. 2, pp. C143–C168, 2015.
- [10] S. Reiter, A. Vogel, I. Heppner, M. Rupp, and G. Wittum, "A massively parallel geometric multigrid solver on hierarchically distributed grids," *Computing and Visualization in Science*, vol. 16, no. 4, pp. 151–164, 2013.
- [11] A. H. Baker, R. D. Falgout, T. V. Kolev, and U. M. Yang, "Scaling hypre's multigrid solvers to 100,000 cores," in *High-Performance Scientific Computing*, M. W. Berry, K. A. Gallivan, E. Gallopoulos, A. Grama, B. Philippe, Y. Saad, and F. Saied, Eds. Springer London, 2012, pp. 261–279.
- [12] S. Badia, A. F. Martín, and J. Principe, "A highly scalable parallel implementation of balancing domain decomposition by constraints," *SIAM Journal on Scientific Computing*, vol. 36, no. 2, pp. C190–C218, 2014.
- [13] O. Ippisch and M. Blatt, "Scalability test of $\mu\phi$ and the Parallel Algebraic Multigrid solver of DUNE-ISTL," in *Jülich Blue Gene/P Extreme Scaling Workshop*, no. FZJ-JSC-IB-2011-02. Jülich Supercomputing Centre, 2011.
- [14] H. Sundar, G. Biros, C. Burstedde, J. Rudi, O. Ghattas, and G. Stadler, "Parallel geometric-algebraic multigrid on unstructured forests of octrees," in *Proceedings of SC12*. ACM/IEEE, 2012.
- [15] C. Burstedde, O. Ghattas, M. Gurnis, E. Tan, T. Tu, G. Stadler, L. C. Wilcox, and S. Zhong, "Scalable adaptive mantle convection simulation on petascale supercomputers," in *Proceedings of SC08*. ACM/IEEE, 2008.
- [16] M. Kronbichler, T. Heister, and W. Bangerth, "High accuracy mantle convection simulation through modern numerical methods," *Geophysical Journal International*, vol. 191, no. 1, pp. 12–29, 2012.
- [17] G. Ranalli, *Rheology of the Earth*. Springer, 1995.
- [18] T. Isaac, C. Burstedde, L. C. Wilcox, and O. Ghattas, "Recursive algorithms for distributed forests of octrees," *SIAM Journal on Scientific Computing (to appear)*, 2015, <http://arxiv.org/abs/1406.0089>.
- [19] M. O. Deville, P. F. Fischer, and E. H. Mund, *High-Order Methods for Incompressible Fluid Flow*, ser. Cambridge Monographs on Applied and Computational Mathematics. Cambridge, UK: Cambridge University Press, 2002, vol. 9.
- [20] M. Benzi, G. H. Golub, and J. Liesen, "Numerical solution of saddle point problems," *Acta Numerica*, vol. 14, pp. 1–137, 2005.
- [21] H. Elman, V. Howle, J. Shadid, R. Shuttleworth, and R. Tuminaro, "Block preconditioners based on approximate commutators," *SIAM Journal on Scientific Computing*, vol. 27, no. 5, pp. 1651–1668, 2006.
- [22] D. A. May and L. Moresi, "Preconditioned iterative methods for Stokes flow problems arising in computational geodynamics," *Physics of the Earth and Planetary Interiors*, vol. 171, pp. 33–47, 2008.
- [23] D. Rossinelli, B. Hejazialhosseini, P. Hadjidoukas, C. Bekas, A. Curioni, A. Bertsch, S. Futral, S. J. Schmidt, N. A. Adams, and P. Koumoutsakos, "11 pflop/s simulations of cloud cavitation collapse," in *Proceedings of SC13*. ACM/IEEE, 2013.
- [24] V. Karakasis, T. Gkountouvas, K. Kourtis, G. I. Goumas, and N. Koziris, "An extended compression format for the optimization of sparse matrix-vector multiplication," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 10, pp. 1930–1940, 2013.
- [25] J. Milano and P. Lembke, "IBM system Blue Gene solution: Blue Gene/Q hardware overview and installation planning," IBM, Tech. Rep. SG24-7872-01, May 2013.
- [26] K. Yoshii, K. Iskra, R. Gupta, P. Beckman, V. Vishwanath, C. Yu, and S. Coghlan, "Evaluating power-monitoring capabilities on IBM Blue Gene/P and Blue Gene/Q," in *Proc. of the IEEE Int. Conf. Cluster Computing*, Beijing, 2012, pp. 36–44.
- [27] U.S. Energy Information Administration (EIA), "Electric power monthly, with data from january 2015," U.S. Department of Energy, Tech. Rep., March 2015. [Online]. Available: <http://www.eia.gov/electricity/monthly/pdf/epm.pdf>
- [28] S. Zhong and M. Gurnis, "Controls on trench topography from dynamic models of subducted slabs," *J. Geophys. Res.*, vol. 99, pp. 15 683–15 695, 1994.
- [29] T. R. A. Song and M. Simons, "Large trench-parallel gravity variations predict seismogenic behavior in subduction zones," *Science*, vol. 301, pp. 630–633, 2003.
- [30] V. Ratnaswamy, G. Stadler, and M. Gurnis, "Adjoint-based estimation of plate coupling in a non-linear mantle flow model: theory and examples," *Geophysical Journal International*, vol. 202, no. 2, pp. 768–786, 2015.
- [31] L. Ruff and H. Kanamori, "Seismic coupling and uncoupling at subduction zones," *Tectonophysics*, vol. 99, no. 2, pp. 99–117, 1983.
- [32] C. H. Scholz and J. Campos, "The seismic coupling of subduction zones revisited," *Journal of Geophysical Research: Solid Earth*, vol. 117, no. B5, 2012.
- [33] D. Forsyth and S. Uyeda, "On the relative importance of the driving forces of plate motion," *Geophysical Journal International*, vol. 43, no. 1, pp. 163–200, 1975.
- [34] B. Hager and R. O'Connell, "A simple global model of plate dynamics and mantle convection," *J. Geophys. Res.*, vol. 86, pp. 4843–4867, 1981.