



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Extending the applicability of multigrid methods

J. Brannick, M. Brezina, R. Falgout, T. Manteuffel, S.
McCormick, J. Ruge, B. Sheehan, J. Xu, L. Zikatanov

September 27, 2006

SciDAC 2006
Denver, CO, United States
June 25, 2006 through June 29, 2006

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

Extending the applicability of multigrid methods

J Brannick¹, M Brezina², R Falgout², T Manteuffel¹, S McCormick¹,
J Ruge¹, B Sheehan², J Xu³, L Zikatanov³

¹ Department of Applied Mathematics, University of Colorado, Boulder, CO, 80309

² Center for Applied and Scientific Computing, Lawrence Livermore National Laboratory,
Livermore, CA 94551

³ Department of Mathematics, The Pennsylvania State University, University Park, PA, 16802

E-mail: brannick1@llnl.gov, mbrezina@math.cudenver.edu, rfalgout@llnl.gov,
tmanteuf@colorado.edu, stevem@colorado.edu, John.Ruge@colorado.edu,
brendan.sheehan@colorado.edu, xu@psu.edu, ludmil@psu.edu

Abstract. Multigrid methods are ideal for solving the increasingly large-scale problems that arise in numerical simulations of physical phenomena because of their potential for computational costs and memory requirements that scale linearly with the degrees of freedom. Unfortunately, they have been historically limited by their applicability to elliptic-type problems and the need for special handling in their implementation. In this paper, we present an overview of several recent theoretical and algorithmic advances made by the TOPS multigrid partners and their collaborators in extending applicability of multigrid methods. Specific examples that are presented include quantum chromodynamics, radiation transport, and electromagnetics.

1. Introduction

The demands of modern numerical simulations in the sciences continue to tax current computational capabilities. Numerical simulation efforts have relied on and benefitted from dramatic improvements in high-performance computer architectures. However, to meet these ever-expanding demands, hardware improvements must be coupled with advances in numerical methods. The left graph in Figure 1, which is from the final report of the *Second DOE Workshop on Multiscale Problems*, Broomfield, Colorado, July 20-22, 2004, shows that advances in algorithms have dramatically improved solver times for large-scale discretized partial differential equations (PDEs). Moreover, comparison with the right graph in Figure 1 from the same report shows that algorithm improvement has kept pace with hardware advances. A substantial part this increase in enabling methodology is due to developments in the multigrid (MG) field.

Multigrid methods entered the modern era of computation in 1977 with Brandt's seminal paper [3], which contained many basic processes used in current MG solvers. Because of its potential for providing solvers that scale linearly with the number of degrees of freedom and number of processors, multigrid has since become the subject of increasing interest and research. However, adoption of MG solvers in applications was slowed by its being perceived only as an elliptic solver whose implementation requires substantial human effort. Several important theoretical and algorithmic advances in multigrid research arose in attempts to address these limitations. As described in the following sections, these developments are beginning to enable the extension of multigrid solvers to a much wider class of applications, many of which are far from the elliptic systems on which the original methodology was based. The purpose of this




Figure 1. Approximate PDE solvers and hardware speed-up (courtesy S. Plimpton and J. Shadid)

paper is to highlight the advances made by TOPS in extending MG applicability. In addition, we point out new areas of research in the MG field where these promising new developments could have a substantial impact.

The remaining sections are organized as follows. Section 2 contains a brief review of: the basic components of algebraic multigrid (AMG); the *element-based* (AMGe) and *adaptive* algebraic multigrid (α AMG) methodologies; and several theoretical results related to such solvers. In Section 3, we present numerical experiments demonstrating that our adaptive smooth aggregation solver is applicable to the linear systems arising in quantum chromodynamics (QCD). Then, in Section 4, we discuss our recent success [15] in applying Multigrid to the systems arising in radiation transport. The success of our approach is due to a robust multi-cell block-Jacobi smoother that allows for standard MG coarsening. In Section 5, we briefly introduce several new approaches for designing optimal AMG-based methods for Maxwell's equations, and Section 6 contains concluding remarks.

2. Background

2.1. Basic AMG components

Multigrid methods are called *scalable* or *optimal* because of their potential to solve $N \times N$ linear systems with only $O(N)$ work. Scalability makes it possible to solve ever larger problems on proportionally larger parallel machines in constant time. Multigrid methods achieve this optimality by employing two complementary processes: *smoothing* and *coarse-grid correction*. In the classical setting of scalar elliptic problems, the smoother (or relaxation method) is a simple iterative method like Gauss-Seidel that is effective at reducing high-frequency error. The remaining low-frequency error is then accurately represented and efficiently eliminated on coarser grids via the coarse-grid correction step. However, applying this simple multigrid idea often involves considerable algorithmic research. One must decide which iterative method to use as a smoother, how to coarsen the problem, and how to transfer information between the grids. What is compelling is that a properly designed multigrid solver can be algorithmically scalable in that it converges uniformly with computational cost that depends only linearly on the problem size.

Multigrid methods must generally exploit the character of the near null space of the operator, that is, vectors, x , that are nearly invisible under the action of the operator A : $Ax \approx 0$. For scalar elliptic PDEs, these vectors are geometrically smooth (low-frequency), but, for many problems

of interest the near null space may be huge (dimension $O(N)$) and may contain geometrically oscillatory vectors, making the task of designing an appropriate multigrid solver considerably more challenging. One family of approaches that are well-suited for addressing these challenges is the class of algebraic multigrid (AMG) solvers.

AMG [4, 34] and its important variant smoothed aggregation (SA; [35]) are MG methods that assume little or no information about the underlying physical structures and, as such, are ideally suited for solving unstructured grid problems. AMG has come to describe a whole class of algorithms that use algebraic information in the matrix to construct the basic multigrid components. The AMG framework usually assumes use of a simple pointwise relaxation method and attempts to correct the algebraically smooth error that remains after relaxation by suitable choice of coarse grids, intergrid transfer operators, and coarse-grid equations. To do this, AMG must be able to efficiently characterize this algebraically smooth error. The classical AMG method of Brandt, McCormick, and Ruge [7] uses a characterization based on properties of M-matrices, namely, that the large off-diagonal entries in the system matrix indicate the direction in which smooth error varies slowly. For the details of this method we refer the reader to [34]. Although the algorithm works remarkably well for a variety of problems, including those that have no M-matrix or elliptic character, the M-matrix assumption nevertheless limits its general applicability.

2.2. Element-based AMG methods

The *element* AMG approach (AMGe [11]) and its variants *element-free* AMGe [23], *spectral* AMGe [16], and *spectral agglomerate* AMGe [17] were developed to improve AMG robustness for finite element problems. AMGe differs from standard AMG by requiring access to individual element stiffness matrices, which are used to construct effective interpolation operators. AMGe uses the multigrid heuristic that interpolation should reproduce an eigenmode with tolerance proportional to the associated eigenvalues: an eigenmode with large eigenvalue need not be interpolated well, while an eigenmode with small eigenvalue must. Though this heuristic gives good guidance for constructing multigrid algorithms, it is impractical because it involves knowledge of the global spectrum of the operator. Instead, AMGe localizes this heuristic to the spectra of small local operators, obtained through summing together the finite element matrices in a small neighborhood. Although AMGe is robust for difficult non-grid-aligned anisotropic diffusion and thin-body elasticity problems, it suffers from generally expensive setup costs because coarse-element matrices are required on all levels. In addition, the coarse problems are generated under the assumption that the error not effectively treated by the smoother varies slowly in algebraic neighborhoods.

Spectral AMGe was developed to handle problems in which this algebraically smooth error is locally oscillatory. This spectral method differs from AMGe in that the degrees of freedom do not make up a subset of the fine degrees of freedom; rather, they are coefficients in the expansion of algebraically smooth vectors in terms of a coarse smooth basis. The coarse basis is constructed locally in terms of eigenvectors of local (small) matrices, and requires an agglomeration procedure similar to the one mentioned above. Tests indicate that the new method with standard smoothing (like Gauss-Seidel) is very efficient and extremely robust. Particularly promising is the fact that its performance seems insensitive to the way the agglomerates are constructed. The main limitation of this method is its need of the local finite element stiffness matrices to construct coarse problems.

2.3. Adaptive AMG

Although these attempts to develop algebraic multigrid methods that apply to more general problems have led to substantial progress in extending applicability of AMG, they are nevertheless restricted by reliance on some assumption about algebraically smooth error and/or

the origin of the problem. Eliminating the need for these assumptions and thereby expanding multigrid into more applied areas is the aim of our research in developing adaptive AMG [12, 13, 10].

Adaptive AMG is a learning approach that iterates on certain artificial problems to uncover the nature of components that must be more effectively resolved in the multigrid coarsening process. To develop an AMG scheme applied to $Ax = b$ for a given matrix A , it is first developed for the homogeneous problem, $Ax = 0$. At the start, relaxation alone is applied to $Ax = 0$. If it converges well, then no coarsening is needed and relaxation alone is acceptable for solving $Ax = b$. Otherwise, the resulting approximation must be algebraically smooth, so it can be used in any of the AMG schemes to define interpolation and restriction operators. To determine if the resulting AMG scheme is effective for the general case, it too is applied to $Ax = 0$. If this current solver is still slow, then the resulting approximation must be an algebraically smooth error that is significantly different than what has already been used to construct this solver. This approximation can then be used in the AMG process to enhance coarsening so it captures all errors with similar local character. Coarsening processes based on emerging errors in the solution of the homogeneous problem must be carefully designed, and understanding the basic principles becomes tricky on coarser levels, but our experience shows that this adaptive process can produce very robust solvers for a much wider class of applications (most notably for QCD [8]).

Another important technique for adaptive multigrid processes is *compatible relaxation (CR)*, recently proposed by Achi Brandt [5]. In its simplest form, CR is just F-relaxation (relaxation at points that do not correspond to coarse-grid points). The idea is that if CR is fast to converge, then the coarse grid is adequate for eliminating the remaining error; but if it is not, then either additional or more aggressive smoothing can be done, or more points can be added to the coarse grid. One crucial difficulty that CR can address is the somewhat large complexity that AMG exhibits in the context of certain applications in three spatial dimensions. For example, AMG can produce denser coarse-level matrices for certain systems of PDEs. One approach to address this issue is aggressive coarsening, that is, choosing a coarse grid with many fewer equations than standard coarsening would allow. CR can be a powerful tool for producing aggressive coarsening that is accurate enough to maintain scalability. We refer the reader to our work in [5, 20, 9] for the details on CR.

2.4. AMG theory

It is important to note that many of the algorithmic advances in AMG are in large part the result of a deeper theoretical understanding of these methods. Classic multigrid theories such as that in [2] and a more recent theory [37] for subspace correction methods have proved useful for design and analysis of several emerging AMG methods. Many theoretical tools have also specifically been developed for AMG, mostly related to two-level methods. The underlying theory for AMGe was outlined in [11] and is based on the *weak approximation property* that, if satisfied by interpolation, implies uniform two-grid convergence. This approximation property relates the accuracy of interpolation to the spectrum of the system matrix: eigenmodes with small associated eigenvalue must be interpolated well. Other methods based on this theory are the so-called energy minimization methods. The problem is that this theory is limited to simple pointwise smoothers and a particular type of coarse grid. We developed a new two-level theory in [20], motivating the use of CR and allowing for general smoothing processes and coarse grids (e.g., vertex-based, cell-based, and agglomeration-based); thus, encompassing a much broader class of problems and algorithms. The motivation for this work was Maxwell's equations, for which pointwise smoothers are inadequate and non-standard coarse grids are often more appropriate. A more recent development is our new sharp theory [21] that gives necessary and sufficient conditions for two-grid convergence and provides additional insight for

the development of AMG methods. While these new theories should enable existing tools to be used to develop more robust AMG methods, the use thereof in designing AMG algorithms is still in very early stages of development.

3. Quantum chromodynamics

QCD explains how neutrons and protons are bound inside nuclei and also how their constituents, gluons and quarks, interact. Since this quantum field theory describes strongly interacting elementary particles, a perturbation theory analysis of it fails to converge. As a result, the dynamics of these interactions must be studied using lattice gauge theory (LGT). The goal in LGT is to compute observables of the quantum theory of a field, $\phi(x)$, using path integrals:

$$\langle O \rangle = \frac{1}{\mathcal{Z}} \int \mathcal{D}\phi O(\phi) e^{-S(\phi)},$$

where the action, S , accounts for the interactions of the theory, and $\mathcal{Z}$ is a partition function. The gauge theory of QCD describes interactions between quarks and gluons and its action is typically written in terms of the *gauge* and *fermionic* actions: $S_{QCD} = S_g + S_f$. The primary task in QCD simulations is accounting for the fermionic action, $S_f = \bar{\psi} M[U] \psi$, where $M[U] = \gamma_\mu D_\mu + m$ is the Dirac operator, m is the quark mass, γ_μ are the usual Dirac matrices, and D_μ are covariant derivatives. Computing this action requires repeatedly computing the action of the inverse of the discretized Dirac operator.

Various discrete models of this operator have been developed, most notably staggered fermions, Wilson fermions, and overlap fermions. For the sake of brevity we limit our discussion here to Wilson fermions and the resulting Wilson fermion matrix. We mention that the difficulties associated with solving systems involving this operator are representative of those encountered in all discrete QCD models; thus, the promising results obtained by applying adaptive AMG to the Wilson fermion system reported here should carry over to other QCD models. The Wilson fermion matrix defines a nearest neighbor coupling of the fermionic degrees of freedom on a four-dimensional hypercube space-time lattice. The link variables, $U_\mu(x)$, represent the gluon fields and are defined on the link in the positive μ direction originating from lattice site x . These link variables take on values in the gauge group $SU(N_c)$ with $N_c = 3$ representing the number of colors in the theory. Their distribution across the lattice is random with the level of randomness prescribed by a temperature parameter β . Typically, for β small (large) the distribution of the gauge field is more (less) random. Given a configuration, U , the Wilson fermion matrix is

$$M[U]_{xy} = \delta_{xy} - \kappa \sum_{\mu} U_{\mu}(x) (1 - \gamma_{\mu}) \delta_{x+\hat{\mu},y} + U_{\mu}^{\dagger}(x - \hat{\mu}) (1 + \gamma_{\mu}) \delta_{x-\hat{\mu},y},$$

where κ is a function of the quark mass, γ_{μ} are the 4×4 Dirac matrices, and $\hat{\mu}$ are unit vectors in directions $\mu = 1, \dots, 4$. Note that there are a total of 12 unknowns per lattice site.

The three main difficulties encountered when attempting to compute actions of the inverse of the Dirac-Wilson operator are: the condition number of the system matrix grows as the quark mass approaches its physical value, implying that a highly accurate AMG interpolation operator must be constructed to ensure that the weak approximation property is satisfied; the near kernel components are geometrically oscillatory; and these oscillations depend on the randomness of the gauge field, which is itself randomly prescribed. As already mentioned, these are precisely the difficulties that the adaptive AMG methodology is intended to address.

Although attempts in the 1990's to introduce multi-scale algorithms to QCD (e.g., [14]) resulted in substantial theoretical progress, they failed for the most part to produce significant advantages for actual QCD simulations of that day. Our preliminary success in applying adaptive

AMG solvers to QCD systems reported in [8] affirms that the circumstances responsible for this failure have dramatically changed. Therein, results of 2D prototype tests are provided, suggesting that adaptive smoothed aggregation α SA [12] may eliminate the so-called “critical slowing down” that remains the main bottleneck in current state-of-the-art QCD simulations. The advantage of α SA is that its coarsening process can automatically identify the near-kernel components resulting from fluctuations of the underlying media (i.e., gauge fields). In Figure 2, we reproduce results from [8], illustrating that the recently developed α SA solver enables the quark mass to approach its critical value with very little additional computational penalty. We note that these results can be contrasted with earlier efforts in [14] that failed in this limit applied to exactly the same prototype 2D Wilson fermion system, known also as the *Schwinger* model in quantum electrodynamics. Reasons for this success are not yet clear, but form a starting point for promising new avenues of research in this area.




Figure 2. Average number of α SA-PCG iterations needed to reduce relative residual by 10^6 vs. the mass constant. Here, the colored vertical lines, corresponding to various choices of β , indicate the values of the relative mass for which the system matrix becomes nearly singular.

4. Radiation transport

Radiation transport is a key issue in a variety of DOE applications, ranging from modeling core collapse supernovae to inertial confinement fusion. Particle transport covers a broad range of phenomena, including the transport of photons, neutrons, neutrinos, and charged particles, each with unique characteristics. For exposition, the discussion here focuses on thermal radiative transfer, which embodies many of the difficulties encountered in other such applications. The processes of absorption and reemission of photons as they move through different material alter the temperature, T , of the material and, thus, its physical properties. These processes are modeled using the Boltzmann transport equation for photon intensity, $\psi(\mathbf{x}, \Omega, \nu, t)$, as a function of space, $\mathbf{x}$, direction, Ω , energy, ν , and time, t , together with an energy balance equation for

the material:

$$\frac{1}{c} \frac{\partial \psi}{\partial t} + \Omega \cdot \nabla \psi + \sigma_t \psi = \frac{1}{4\pi} \sigma_s \phi + \sigma_a B(\nu, T), \quad (1)$$

$$C_p(T) \frac{\partial T}{\partial t} = \int_{\nu} \sigma_a (\int_{\Omega} \psi - 4\pi B(\nu, T)). \quad (2)$$

Here, $\phi = \int_{\Omega} \psi$ is the angular integrated radiation intensity, c is the speed of light, $\sigma_t(\mathbf{x}, \nu, T)$ is the total macroscopic collision cross-section, $\sigma_s(\mathbf{x}, \nu, T)$ is the macroscopic scattering cross-section, $\sigma_a(\mathbf{x}, \nu, T)$ is the macroscopic absorption cross-section, $C_p(T)$ is the material heat capacity, and $B(\nu, T)$ is the Planck function. To simplify this discussion, scattering has been modeled as monochromatic and isotropic, although in practical applications it is neither. Linearization about a current T , backward-difference approximation in time, and integration over G discrete energy intervals yields the multi-group transport equation

$$\Omega \cdot \nabla \psi_g + \sigma_{t,g}^* \psi_g - \frac{1}{4\pi} \sigma_{s,g} \phi_g = \frac{1}{4\pi} C_g \sum_{k=1}^G \sigma_{a,k} \phi_k + \xi_g, \quad \text{for } g = 1, \dots, G, \quad (3)$$

where $\sigma_{t,g}^* = \sigma_{t,g} + \frac{1}{c\Delta t}$ is the effective cross-section and C_g and ξ_g are now known quantities. This system of equations typically involves heterogeneous material whose cross-sections can differ by several orders of magnitude. In optically thick material, the solution may be diffusive in nature while, in optically thin material, the solution takes on a hyperbolic character.

If the total absorption rate, $f_a = \sum_{k=1}^G \sigma_{a,k} \phi_k$, is known, the above system breaks into separate single-group equations. This motivates the grey transport acceleration scheme (GTA; e.g., [28]) in which the single-group equations are solved and used to update the value of f_a . A grey transport equation involving f_a is used to accelerate the iteration. This grey transport equation can be approximated by a drift-diffusion equation, which can be solved more efficiently if multigrid techniques are available. Similarly, if the group angular integrated intensity, ϕ_g , is also known, each single-group equation breaks into separate hyperbolic systems involving the first two terms in (3). This is the motivation for the diffusion synthetic acceleration scheme (DSA; [1]), in which the hyperbolic system is solved for a discrete set of angles, called discrete ordinates, and used to update the value of ϕ_g . The iteration is accelerated by a diffusion equation involving ϕ_g . Thus, the multi-group transport equation is solved by an inner iteration, DSA, for the single-group equations, and an outer iteration, GTA, for the multi-group equation. This is done at each step of a Newton-like iteration for an implicit time-stepping scheme.

In these solution strategies, multigrid plays its traditional role as an elliptic solver. Both the drift-diffusion of GTA and the diffusion equation in DSA, being elliptic in nature, benefit from multigrid algorithms based on standard multigrid technology. However, research in recent years has focused on designing spatial multigrid methods to directly address the single-group equations, which are not strictly elliptic. To apply spatial multigrid to the single-group equations, a smoothing step must be identified that facilitates spatial coarsening. Early work on multigrid algorithms for this equation in slab geometry employed a block Jacobi relaxation, where the blocks correspond to two-cell pairs on the spatial grid [33]. A similar algorithm was implemented in two spatial dimensions using a block Jacobi based on 4-cell blocks [27]. More recent results show that in multiple spatial dimensions, a single block Jacobi step is not sufficient to accomplish spatial smoothing. To overcome this limitation we developed a multi-pass block-Jacobi smoother, where the spatial blocks are shifted after each pass and the number of passes needed is 2^d , with d denoting the number of spatial dimensions. In [15], a multigrid algorithm employing this shifted block relaxation was tested in two spatial dimensions for the corner balance and discontinuous Galerkin discretizations. The convergence factors of the resulting algorithm are less than 0.1 for a variety of homogeneous and heterogeneous material properties.

Naturally, convergence degrades in the case of a vacuum, where the equation is purely hyperbolic. The vacuum case, however, has also been addressed with a multigrid algorithm based on a least-squares discretization [19].

5. Electromagnetics

The scalable solution of large-scale electromagnetics systems is important in a number of DOE applications, including accelerator design and magnetic confinement fusion (magnetohydrodynamics). Three basic systems resulting from electromagnetics applications are: the *semi*-definite Maxwell's equations or so-called time-domain equations; the *indefinite* Maxwell's system or so-called frequency domain equations; and the Helmholtz equation. All three lead to ill-conditioned discrete systems having huge near null spaces, thus making the development of effective multigrid methods extremely difficult. Unstructured meshes are often required, further increasing the difficulty of the problem and making AMG a natural approach for developing optimal solvers.

Often, a simplified system, known as the eddy current equations, is used in computational models of the time harmonic Maxwell's equations. This system is obtained by neglecting the displacement current and the high frequency speed-of-light time scale electromagnetic waves in a conducting media and using an implicit discretization in time. This results in the following semi-definite system of PDEs for the components of the electric field $\mathbf{E} \in H(\mathbf{curl})$ to be solved on each time step:

$$\mathbf{curl}(\alpha \mathbf{curl} \mathbf{E}) + \beta \mathbf{E} = \mathbf{f}, \quad (4)$$

where $\alpha > 0$, $\beta \geq 0$ and $\mathbf{f}$ are given functions, and $\mathbf{f}$ is divergence free. We note that when nonconducting regions are present, β can be zero and the system becomes singular. One of the important properties of this system is that it has an infinite-dimensional near null space (gradients of H^1 functions). This property clearly limits the set of admissible discretizations that can be used in computing a numerical solution to (4). In fact, it is well known [24] that a consistent discretization of the higher-order term in this system must satisfy certain compatibility conditions, dictated by the infinite-dimensional kernel of the $\mathbf{curl}$ operator.

Commonly used discretizations of the eddy current equations, satisfying the appropriate compatibility conditions, are based on the Nédélec edge finite elements. Indeed, these finite element spaces have a proper representation of the null space of the discretized $\mathbf{curl}$ operator, namely, the space of gradients of piecewise polynomial functions, and, thus, provide a tool for building numerical models of (4). These consistent numerical models naturally preserve the features of the continuous problem and, therefore, the resulting discretized system is often very ill conditioned (or even singular). In addition, an accurate approximation of the electric field requires a huge number of degrees of freedom for which standard solution approaches simply do not work. Standard multigrid methods using pointwise smoothers cannot eliminate a sufficient number of the near null components, because this requires too large of a coarse grid. To resolve this issue, special smoothers must be used to damp the locally supported (oscillatory) null space components of the fine-grid operator, thus allowing for a standard coarse-grid correction. This need for more powerful smoothers is what motivated the theoretical work found in [20, 21]. These more robust smoothers are, however, only viable on uniformly refined grids.

Recent theoretical work in [25] has led to a new AMG preconditioner for the solution of variational problems in $H(\mathbf{curl})$ and $H(div)$. The theoretical foundation of this new approach is the auxiliary space preconditioning framework [36], with an auxiliary space constructed using *regular splittings* of $H(\mathbf{curl})$ and $H(div)$. Although this theory is now only developed for definite equations it should in principle also be applicable to semi-definite and indefinite cases. Using this technique, h -independent preconditioners can be developed by employing any standard preconditioner for scalar elliptic equations, e.g. AMG. For $H(\mathbf{curl})$ equations, the method uses

a standard relaxation scheme and a solver for several (four in case of $H(\mathbf{curl})$ and seven in case of $H(\mathbf{div})$) scalar elliptic equations. The rigorous theoretical analysis in [25] and the extensive numerical experiments in [26] indicate that this approach leads to a robust preconditioner for discretized $H(\mathbf{curl})$ and $H(\mathbf{div})$ systems. We mention that, although this approach relies on geometric information, this is only needed on the finest grid, which is not a serious limitation for most practical applications. Further research must, however, be conducted to alleviate such grid dependence, as well as to extend this method to the semi-definite and indefinite cases.

Recent work [30, 18] has demonstrated that by making direct use of the divergence-free constraint, $\mathbf{div}(\sigma\mathbf{E}) = 0$, not explicitly accounted for in (4), and then casting the problem as a first-order system leads to a continuous operator that does not have an infinite-dimensional near nullspace and, so, standard multigrid methods can be very effective as solvers for this formulation. However, in this setting, care must then be taken in the presence of singularities in the solution of the continuous problem. We have developed a FOSLL* formulation [30] to address these issues.

The Helmholtz equation, $-\nabla^2 E - k^2 E = f$, also appears in many electromagnetics calculations. This equation models the scattering of waves due to an obstacle. The main difficulty in solving the discrete systems resulting from this equation is the near null space comprised of plane waves that are geometrically smooth in some direction and oscillatory in the perpendicular direction. The varying nature of these directions means that, typically, these error components cannot be accurately represented using standard coarse grids. There are two multigrid algorithms for solving the Helmholtz equation that have demonstrated convergence independent of mesh size h and wavenumber k : the wave-ray algorithm of Brandt and Livshits [6, 32] and our first-order system least squares (FOSLS) approach [29]. Both use the computationally expensive approach of computing multiple coarse spaces to accurately approximate the plane waves in the near null space. Nonetheless these results demonstrate the existence of a multigrid solver for such systems. Developing more efficient approaches for such systems provides for another interesting area of future research in the development of multigrid methods.

6. Concluding remarks

Recent theoretical and algorithmic advances made by the TOPS multigrid partners and their collaborators have led to the design of significantly more robust AMG algorithms. The extended applicability of MG methods due to these advances in turn have the potential to allow for more complex and higher fidelity physics for several important SciDAC applications, for example, lattice QCD, radiation transport, and electromagnetism. Indeed, our promising preliminary results obtained by applying adaptive smooth aggregation to the discrete problems arising in lattice QCD [8] suggest that adaptive AMG may provide for an optimal solver for these linear systems. Our successful application of spacial multigrid to the systems in radiation transport [15] provide for yet another promising area of future research in developing MG methods. The recent advances made in developing AMG-based solvers for discretized Maxwell's equations [25, 30] should provide for an optimal solver for this long outstanding open problem. Further development of these algorithms and scalable parallel implementations thereof are examples of the many exciting new areas of research emerging in the MG field.

References

- [1] Alcouffe R 1976 A stable diffusion synthetic acceleration method for neutron transport iterations *Trans. Am. Nucl. Soc.* **23** 203
- [2] Bramble J, Pasciak J, Wang J, and Xu J 1991 Convergence estimates for multigrid algorithms without regularity assumptions *Math. Comp.* **57** (195) 23–45
- [3] Brandt A 1977 Multi-level adaptive solutions to boundary value problems *Math. Comp.* **31** 333–390
- [4] Brandt A 1986 Algebraic multigrid theory: The symmetric case *Appl. Math. Comp.* **19** 23–56

- [5] Brandt A 2000 General highly accurate algebraic coarsening *ETNA* **10** 1–20
- [6] A. Brandt and I. Livshits 1997 *Wave-ray multigrid method for solving standing wave equations*, *ETNA* **6** 162–181.
- [7] Brandt A, McCormick S, and Ruge J 1982 Algebraic multigrid (AMG) for automatic multigrid solutions with application to geodetic computations *Report Inst. for Comp. Studies* Fort Collins CO
- [8] Brannick J, Brezina M, Keyes D, Livne O, Livshits I, MacLachlan S, Manteuffel T, McCormick S, Ruge J, and Zikatanov L Adaptive smoothed aggregation in lattice QCD *Lecture Notes Comp. Sci. Eng.* Springer-Verlag, to appear.
- [9] Brannick J and Falgout R Compatible relaxation and coarsening in AMG, in preparation.
- [10] Brannick J and Zikatanov L, AMG methods based on compatible relaxation and energy minimization, *Lecture Notes Comp. Sci. Eng.*, Springer Verlag, to appear.
- [11] Brezina M, Cleary A, Falgout R, Henson V, Jones J, Manteuffel T, McCormick S, and Ruge J Algebraic multigrid based on element interpolation (AMGe) 2000 *SIAM J. Sci. Comp.* **22** 1570–1592.
- [12] Brezina M, Falgout R, MacLachlan S, Manteuffel T, McCormick S, and Ruge J 2004 Adaptive smoothed aggregation (α SA) *SIAM J. Sci. Comp.* **25** (6) 1896–192
- [13] M. Brezina, R. Falgout, S. MacLachlan, T. Manteuffel, S. McCormick, and J. Ruge 2006 Adaptive algebraic multigrid *SIAM J. Sci. Comp.* **26** 1261–1286.
- [14] Brower R, Rebbi C, and Vicari E 1991 Projective multigrid method for propagators in lattice gauge theory *Phys. Rev.* **D43** 1965–1973.
- [15] Chang B, Manteuffel T, McCormick S, Ruge J, and Sheehan B 2006 Spatial multigrid for isotropic neutron transport *SISC*, submitted.
- [16] Chartier T et. al. 2003 Spectral AMGe (ρ AMGe), *SIAM J. Sci. Comp.* **25** 1–26.
- [17] Chartier T et. al. 2005 Spectral agglomerate AMGe, *Lecture Notes Comp. Sci. Eng.* Springer-Verlag, to appear.
- [18] Costable M and Dauge M 2000 *Numer. Math.* **93** 239–277.
- [19] De Sterck H, Manteuffel T, McCormick S, and Olson L 2004 Least-squares finite element methods and algebraic multigrid solvers for linear hyperbolic PDEs *SIAM J. Sci. Comp.* **26** 31–54.
- [20] Falgout R and Vassilevski P 2004 On generalizing the algebraic multigrid framework *SIAM J. Num. Anal.* **42** 1669–1693.
- [21] Falgout R, Vassilevski P, and Zikatanov L 2005 On two-grid convergence estimates *NLAA* **12** 471–494.
- [22] *hypr*: High performance preconditioners, <http://www.llnl.gov/CASC/hypr/>
- [23] Henson V and Vassilevski P 2001 Element-free AMGe: general algorithms for computing interpolation weights AMG *SIAM J. Sci. Comput.* **23** 629–650.
- [24] Hiptmair R 1999 A Multigrid method for Maxwell’s equations *SIAM J. Num. Anal.* **36** 204–225.
- [25] Hiptmair R and Xu J Preconditioning $H(\text{curl})$ and $H(\text{div})$ systems using Poisson solvers, in preparation.
- [26] Kolev T, Vassilevski P 2006 Some experience with a H^1 -based auxiliary space AMG for $H(\text{curl})$ problems *Tech. Report LLNL UCRL-TR-221184*.
- [27] Lansrud B and Adams M 2005 A spatial multigrid iterative method for two-dimensional discrete-ordinates transport problems *Proc. Math. and Comp.*, to appear.
- [28] Larsen E 1988 A grey transport acceleration method for time dependent radiative transfer problems *J. Comput. Phys.* **78**, 459.
- [29] Lee B, Manteuffel T, McCormick S, and Ruge J 2000 First-order system least squares for Helmholtz equations, *SIAM J. Sci. Comp.* **21** (5) 1927–1949.
- [30] Lee E and Manteuffel T 2006 FOSLL* Method for the Eddy Current problem with three dimensional edge singularities, in preparation.
- [31] Lee Y, Wu J, Xu J, and Zikatanov L Subspace correction methods for nearly singular problems, in preparation.
- [32] Livshits I 2004 An algebraic multigrid wave-ray algorithm to solve eigenvalue problems for the Helmholtz operator *NLAA* **11** 229–239.
- [33] Manteuffel T, McCormick S, Morel J, Oliveira S, and Yang G 1994 A Parallel multigrid method for isotropic transport equations, *SIAM J. Sci. Comp.* **15** 474–493.
- [34] Ruge J and Stüben K 1987 Algebraic multigrid (AMG) Multigrid Methods S. McCormick ed. **3** *Frontiers in Applied Mathematics* SIAM Philadelphia PA 73–130.
- [35] Vaněk P, Mandel J, and Brezina M 1996 Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems *Computing* **56** 179–196.
- [36] Xu J 1996 The auxiliary space method and optimal multigrid preconditioning techniques for unstructured meshes *Computing* **56** (3) 215–235.
- [37] Xu J and Zikatanov L 2002 The method of alternating projections and the method of subspace corrections in Hilbert space *J. Amer. Math. Soc.* **15** 573–597.