Research Statement

Scott MacLachlan

scott.maclachlan@gmail.com
http://ta.twi.tudelft.nl/users/maclachl

Delft Institute of Applied Mathematics
Delft University of Technology
Mekelweg 4, Room 07-270
2628 CD Delft
The Netherlands
Office: +31 15 278 3634

Field of study: My interests lie in the fields of scientific computation and computational applied mathematics, with emphasis on multiscale numerical analysis and efficient numerical solution of partial differential equations. My work is motivated by real-world applications, such as fluid flow through porous media, elastic deformation in biological and geophysical settings, quantum chromodynamics, and electromagnetic fields in geophysical media or circuit design. The common theme is characterizing the properties of a mathematical model that are responsible for slow convergence of standard techniques, constructing reduced auxiliary models that appropriately capture these dynamics, and using these models to effectively accelerate convergence. These auxiliary models may also be used in a multiscale modelling framework to achieve improved accuracy relative to computational cost for reduced-order discretizations.

Past and Current Research: My research focuses on the use of coarse-scale models to improve the performance of numerical approaches to solving differential equations. Robust algorithms that appropriately address the inherent heterogeneity found in mathematical models of many interesting real-world processes are required for efficient and relevant simulation. In order to improve existing solver technology, we seek to identify the properties of a given model that cause difficulties in classical approaches and, then, to exploit this knowledge in constructing reduced models that address these barriers to efficiency. From this principle, multigrid methods provide *optimal* solution of many linear systems (that is, in time linearly proportional to the number of discrete degrees of freedom) by combining inexpensive local processing (such as the Jacobi or Gauss-Seidel iterations) with a carefully chosen coarse-scale correction process. Algebraic multigrid (AMG) methods are based on the same principle, but attempt to automatically construct such an optimal solver with few (or no) direct assumptions on the underlying mathematical model but, rather, by making assumptions on the spectral properties of the linear system. This generalization results in AMG often being the solver of choice for finite element discretizations on irregular domains with unstructured meshes.

My Ph.D. research was motivated by a desire to improve the robustness of algebraic multigrid algorithms. Multigrid solvers (AMG, in particular) are known to be extremely effective for the solution of standard discretizations of many elliptic operators, whose spectral properties are relatively easy to characterize. For non-standard models, such as those with materials that vary on a fine spatial scale, effective characterization of the needed properties may be quite difficult. Thus, the main focus of my thesis was the development of a new family of *adaptive* algebraic multigrid methods that do not rely on the classical assumptions on the spectral properties of the linear system. The resulting methods show typical multigrid convergence for scalar, elliptic second-order PDEs and for systems of PDEs, such as linear elasticity, even when no spectral information is given to the solver. Analyzing the simpler case of only two grid levels, we proved the convergence of the adaptive process, as well as detailing a class of problems for which these methods are expected to perform well. This research led to four journal publications, including one selected to appear in the SIAM Review.

This research led to the first successful application of multigrid-type solvers in the field of lattice gauge theory, a discrete approach to the theory of quantum chromodynamics (QCD), describing the basic

building blocks of matter. Within a Monte-Carlo QCD simulation, a series of linear systems must be solved, whose entries are dependent on a gauge field drawn from a known probability distribution. Important properties of these linear systems change with the gauge field, making a static choice of preconditioned Krylov technique for the entire Monte-Carlo simulation ineffective. Classical preconditioners cannot overcome the so-called *critical slowing down* that results when the parameters of the model are set to physically relevant values. The adaptive smoothed aggregation multigrid approach, however, significantly reduces the number of iterations needed to solve each linear system for a reduced, two-dimensional lattice model.

While AMG is often viewed as a black-box approach for solving linear systems, it is based on variational principles that provide useful insight into the finite-element models to which it is often applied. Numerical homogenization, or upscaling, techniques (also known as reduced-order modelling) rely on a consistent reduction of a PDE model from a given (spatial) scale to a coarser scale that is more tractable for computation. To maintain fine-scale accuracy, in terms of a given approximation property or more-general measure, an appropriate basis for a lower-dimensional space must be chosen. Within AMG, operators on discrete coarser scales are formed by a projection of the fine-scale operator onto the space on which localized fine-scale processing is ineffective. Such a space, however, typically corresponds to that which dominates important approximation properties. That is, the variational AMG process implicitly defines a basis for an appropriate coarse-scale space and explicitly computes the projection of the fine-scale operator onto this space. For the Darcy law model of saturated, single-phase flow through a porous media, we have demonstrated that important features of the fine-scale solution are preserved in the solutions of the coarse-scale models created by the AMG process. Such simulations dominate the computational costs of reservoir simulations (for both oil production and environmental reasons); efficient use of these multiscale models allows for current simulations to be performed in a fraction of their current time and enables simulations based on higher-resolution data to be performed as well.

Current efforts, in collaboration with Los Alamos National Laboratory, are aimed at extending both this multilevel upscaling technique and our understanding of it. Contrary to many existing *multiscale* techniques, the AMG framework provides a complete hierarchy of consistent models across many scales, which can be used to automatically select the level of accuracy in the coarse-scale solution. Such a multiscale representation is also quite useful for nonlinear models, where the dominating behaviour may occur on vastly different scales in different parts of the physical domain. Ongoing theoretical analysis focuses on the relationship between the AMG coarsening process and recent advances in generalized and multiscale finite element methods.

In my earlier postdoctoral research, we considered the important practical question of choosing the coarse-grid degrees of freedom within the AMG algorithm. As long as the coarse-scale model is of sufficient dimension, the subspace corresponding to errors that are slow to be resolved by localized fine-scale processing can always be represented. From a practical standpoint, however, compactly supported basis functions for this subspace are necessary to maintain the optimal efficiency of the multigrid process. This localization is only possible if a set of good seed nodes for these basis functions is properly identified. The selection of these nodes is often based on heuristics derived for the matrices arising in isotropic diffusion processes; for more difficult problems, these heuristics may lead to poor choices of coarse-scale nodes, resulting in either a poor-quality correction or one that is too expensive to calculate. Our work identifies theoretical properties of good coarse-scale nodes for any symmetric and definite operator. At the University of Minnesota, we devised a coarsening procedure that, coupled with standard choices of inter-scale transfer operators, guarantees optimal two-level performance and, in practice, delivers optimal or near-optimal multilevel performance for many symmetric and non-symmetric problems.

My current research, as a Marie Curie fellow at the Delft University of Technology, is on the development of multigrid techniques for PDEs with complex-valued coefficients. This project is aimed at developing optimally efficient solvers for applications such as frequency-domain wave propagation, electromagnetics, and QCD, in the presence of heterogeneity. Previous work by the Delft group produced efficient geometric multigrid solvers for the Helmholtz equation; we are investigating further improvements to efficiency that are possible using AMG ideas for the frequency-domain acoustic wave equation (in collaboration with Shell), as well as for applications that naturally lead to unstructured meshes, such as circuit-board design and fabrication. A motivating theoretical framework is provided by the tools of local Fourier analysis that, for the homogeneous case, give sharp bounds on the expected AMG performance. Extending my previous work on multigrid solvers for (real-valued) heterogeneous applications, in combination with this analysis, we aim to develop algebraic multigrid tools for complex-valued problems that are as powerful as those available for real-valued problems.

Future Directions: The ever-increasing power of computational hardware is still not sufficient to meet the simulation needs of application scientists; advanced algorithmic tools are also needed to enable simulations of physically relevant phenomena in interesting parameter regimes. Driving factors in these needs are the inherent heterogeneity of the world around us as well as the multiscale nature of our mathematical models.

Recent experiments suggest that the adaptive AMG framework is a powerful tool for problems where the coefficients of the model depend on random parameters. The theory of quantum chromodynamics is an attempt to understand matter at scales where validating experiments are often not practical; effective algorithms would allow simulation to take some of the role of expensive experiments. Uncertainty in data also leads to models with random parameters in other applications, again requiring adaptive solution techniques to achieve optimal efficiency for each realization. Further development of the adaptive multigrid framework will address the effects of stochastic coefficients on the choices made in the multigrid algorithm and consider how to best make these choices, based on known properties of the random parameters.

Understanding the dynamics of industrial and environmental processes requires efficient and accurate simulation tools. The flow of multiphase reacting fluids in porous media is an important problem in oil-reservoir management, modelling environmental contamination, and developing safe storage and remediation plans for toxic waste. Modelled by a system of nonlinearly coupled PDEs, these applications quickly lead to complicated dynamics across many space and time scales. A flexible model-order reduction framework, such as the AMG-based multilevel upscaling approach described above, will capture the evolving reaction dynamics along boundaries between fluid phases. Through the use of error indicators that naturally arise in the generalization of multigrid to nonlinear problems, the basis-function analogue of automated grid refinement and de-refinement can maintain optimal approximation properties with a minimum number of degrees of freedom. While there are many challenges in realizing such a simulation, the principles on which it is based are well-understood for simpler models. Addressing the mathematical and computational challenges in a systematic way will lead to tools that offer optimal accuracy measured against the computational cost.

The multidisciplinary nature of applied mathematics requires understanding of physical problems, mathematical models, and computational techniques. I aim to develop and understand computational tools that address the simulation needs of application scientists. I am eager to explore new fields of research, and new applications for my past and current work; my experiences in numerical simulation and mathematical analysis lead to natural collaboration with industrial and academic partners. Following this applications-driven approach, I hope to develop a research program blending mathematical analysis and computational science that improves our understanding of the world around us.