From *SIAM News*, Volume 37, Number 2, March 2004

Application Scientists Join Algorithm Designers at Workshop on Large-Scale Nonlinear Problems

By Homer Walker and Carol S. Woodward

Research advances in computational science have made simulation a truly equal partner with theory and experiment in understanding physical phenomena. Decision makers in government and industry increasingly rely on simulations as they make choices of ever-greater impact. As a result, computational scientists are being called on to solve problems with more complex and coupled physics than ever before. In such applications as fluid dynamics, fusion, electronics, groundwater flow, astrophysics, and combustion, coupled models are replacing studies previously limited to individual effects. In addition, parallel computers with large storage capacities have paved the way for high-resolution simulations on large-scale domains. The growth in the complexity and size of models, coupled with the advent of more powerful machines, has led to a continued and increasing demand for effective algorithms for solving large-scale systems of nonlinear equations.

The Workshop on Solution Methods for Large-Scale Nonlinear Problems, held August 6–8, 2003, in Livermore, California, brought together many of the most active researchers on both the algorithmic and the applications sides of the area. The Center for Applied Scientific Computing and the Institute for Scientific Computing Research at Lawrence Livermore National Laboratory hosted the workshop. Major themes included Newton–Krylov methods, preconditioning techniques, operator-split and fully implicit schemes, continuation methods, and PDE-constrained optimization. Among the applications addressed were fluid dynamics, astrophysics, magnetohydrodynamics, porous media flows, and radiation diffusion. The workshop was a follow-on to similar workshops on large-scale nonlinear problems held at Utah State University (1989, 1995) and in Pleasanton, California (2000).

Workshop attendees came from academia (17), government laboratories (25), and industry (2) and were mainly from the United States. Almost a quarter of the participants were graduate students or postdocs (see sidebar for one student's comments on the workshop), and an additional four or five participants were in their first three years past the postdoc—a much larger proportion of young researchers than at the earlier workshops. This high percentage of young researchers in the field reflects growing interest and opportunities in this dynamic area.

Newton–Krylov Methods: "Robustification" and Preconditioning

From the presentations, it was clear that Newton–Krylov methods are still the workhorse methods in the field. In these algorithms, Newton's method is combined with preconditioned Krylov solvers to produce approximate solutions of the linear Jacobian systems. Efficient schemes for achieving fast convergence, up to the quadratic convergence of Newton's method, have been realized in a number of applications, leaving "robustification" as the main focus of current research. Accordingly, a number of workshop speakers concentrated on globalizations and continuation techniques combined with Newton–Krylov methods.

The most commonly applied globalization has been *linesearch* (*backtracking, damping*) globalization, in which each step direction is that of the approximate solution of the Jacobian system and the step length is chosen to give desirable progress toward a solution. *Trust-region methods* offer an alternative approach, in which each step is taken with the goal of optimally reducing the norm of the local linear model within a "region of trust" of the model around the current approximate solution. Trust-region implementations of Newton's method combined with the conjugate gradient method (known as *truncated Newton methods*) have been applied to large-scale optimization problems for some time. In contrast, trust-region globalizations of more general Newton–Krylov methods were used in the past only rarely; recently, however, they have received renewed attention as new issues and possibilities have come to light. Pseudo-transient continuation methods, used in computational fluid dynamics for a number of years, are now starting to be invoked in other application areas, such as groundwater flow. Similarly, natural-parameter continuation methods, used for some time in certain applications, are combining with Newton's method in a broader range of applications. Systematic comparisons of all these robustification techniques have begun, and further understanding of which method is best for a particular class of problems is a subject of active research.

Some workshop speakers addressed recent advances in the use of Newton's method. One such advance is the development of theory for extensions of the method to certain classes of nonsmooth nonlinear functions. Using the concept of generalized derivatives, researchers have shown convergence of Newton's method for these functions, both with and without pseudo-transient continuation. Another advance is in the use of automatic differentiation for efficient generation of accurate Jacobian–vector products within a Newton–Krylov method; this approach has been shown to have computational speed advantages over finite-difference approximations for fluid dynamics problems. Additionally, initial results were given for the application of a two-grid technique that transfers nonlinearities to a coarser-resolution version of the original problem. Lastly, speakers reported new developments in globalized tensor–Krylov methods, which can be regarded as extensions of Newton–Krylov methods that incorporate limited second-order information. Recently developed globalizations, in combination with certain extensions of GMRES, provide robustness and efficiency while allowing the superlinear convergence associated with direct tensor methods on singular and ill-conditioned problems.

Preconditioning for both nonlinear and linear problems continues to be an important area of research. Speakers outlined further advances in preconditioning the nonlinear problem via a nonlinear additive Schwarz approach, including three-dimensional and parallel results on computational fluid dynamics problems. Preconditioning the linear Jacobian systems continues to be a primary requirement for the success of Newton–Krylov methods on large-scale problems. In a physics-based preconditioning approach that has had significant success in many application areas, the preconditioner is formed by either splitting or lagging the physics in the problem in order to capture the most important relationships. In a reaction–diffusion problem, for example, the preconditioning step can be done by first solving the diffusion portion of the problem and then solving the reaction part. Often, this approach can reuse simulation code developed previously for an explicit or operator-split formulation. Multigrid preconditioners continue to deliver scalability for very large problems and are easily embedded in a physics-based preconditioning approach. Lastly, sparse approximate inverse preconditioners have also been used effectively for large-scale, highly coupled problems.

Applications

Applications scientists at the workshop reported on progress resulting from the incorporation of some of these algorithmic advances. In some areas of fusion simulation, application of Newton–Krylov methods has enabled a move to implicit formulations that use larger time steps that are not limited by the fastest wave speed in the system. Speakers also described extensions of these approaches to the more complicated case of implicit formulations of fusion problems on adaptive meshes.

Groundwater and geomechanics codes are being coupled, as are shallow-water and groundwater simulators. Although these couplings are not all yet fully implicit, nonlinear couplings between the relevant effects can now be studied. In the case of fully implicit formulations of shallow-water problems, improved preconditioning and nonlinear solvers allow simulation of hurricanes with much larger time steps than possible previously. In other applications, preconditioned Newton–Krylov methods have led to progress in the simulation of combustion chemistry and in the modeling of phase transitions in smart materials. In astrophysics, these methods have also made possible implicit formulations with larger time steps in simulations of multigroup flux-limited diffusion of neutrinos within core-collapse supernovae simulations.

Discussions and Posters

Each day of the workshop ended with a moderated discussion. The first day's topics were robustness and failure of nonlinear solvers. A review of types of failure—including divergence or stagnation of iterates, convergence to a local norm minimizer that is not a solution, failure of the linear solver, and convergence to a "wrong" solution—was followed by a discussion of causes, symptoms, and possible remedies. Participants also considered the general question of how to construct nonlinear solver algorithms that either prevent failure or terminate with useful diagnostic information. Concluding the session, participants described their experiences with several globalization methods (including continuation) and assessed the relative effectiveness of the methods in practice.

Tolerances, stopping criteria, and related accuracy/efficiency considerations for nonlinear solvers were the main themes of the second day's discussion. The initial focus was on fully implicit methods for time-dependent problems, with a review of their merits and disadvantages relative to explicit and semi-implicit methods. Participants then addressed the choice of stopping tolerances for nonlinear residuals associated with the implicit equations, generally agreeing that a relative stopping criterion, such as 10^{-5}, works well in most instances, although some prefer absolute tolerances

Newton–Krylov Methods Come Out of the Textbooks

Actively interested young people, Homer Walker and Carol Woodward write in the accompanying article, are surely a mark of a dynamic field. Among the many young people at the workshop they describe was Joseph Simonis, a graduate student at Worcester Polytechnic Institute, soon to begin his dissertation research in the area of numerical methods for nonlinear PDEs. At the workshop, Simonis presented a poster describing work done during a summer internship at Sandia National Laboratories. Contacted after the workshop, he agreed to provide SIAM News *readers with a perspective from a person just entering the field.*

My poster presented the results of a numerical study done at Sandia National Laboratories in collaboration with John Shadid, Roger Pawlowski, and my adviser, Homer Walker. Our aim was to evaluate the relative merits of several globalized Newton–GMRES methods on large-scale 2D and 3D problems involving the steady-state Navier–Stokes equations.

Coming into the workshop, I knew about the Newton–Krylov methods in an academic sense. They are powerful tools, used to solve lots of interesting problems, so of course I "knew" them, but until the workshop I don't think I really comprehended the breadth of the problems these methods can solve. I found that people are using them for a very wide spectrum of problems, seemingly unrelated but all having, at their core, the same numerical methods.

It's something like seeing the Grand Canyon for the first time. You know it's going to be large, you can read about its dimensions, but until you see it for yourself, you don't completely believe that its really big. It's much the same with these numerical methods. Everyone says they're useful—they rattle off a bunch of problems—and I've used them in classes to solve some problems. But at the workshop, I got to see countless problems for which the methods were successfully used.

The workshop gave me new motivation for my studies. Seeing results and evidence that what I'm studying is useful makes the studying go so much easier.

The presence of other students and postdocs definitely made the experience more enjoyable. For one thing, it gave me the chance to interact with other people of my age and mathematical experience; I got to talk to other students who are struggling with the same problems and have similar levels of understanding. I also learned about programs at other schools, and what is expected of others from their advisers.

I thoroughly enjoyed my national lab experience. The Sandia internship was very valuable in that I've learned about lots of really interesting problems and have met many other students from all over the country working on many interesting problems. I've also found ways to answer the age-old question posed by so many of my relatives: "What do you do with a degree in math?"

in some applications. Turning to choice of the difference step in finite-difference approximations of Jacobian–vector products in "matrix-free" Newton–Krylov implementations, participants agreed that difference-step formulas should include inner products as well as norms of the relevant vectors. On the choice of relative residual norm tolerances for the approximate solution of Jacobian systems (the *forcing terms*), participants reported a variety of experiences in particular applications, ranging from adaptive formulations to relatively large (10^{-1}) and small (10^{-4}) constant choices.

The third day's discussion centered on concerns for nonlinear solver software writers and algorithm designers if they are to take advantage of upcoming architectures. Two main consensus items emerged: careful use of data, so that cache hierarchies can be exploited and data reuse maximized, and fault tolerance in implementations. With the large numbers of processors on platforms like the new ASCI machines and Blue Gene/L, it is likely that some number of processors will fail and drop out of a simulation unexpectedly. Thus, fault-tolerant and robust solver implementations will be required for the successful use of such machines.

A poster session held the evening of the second day highlighted the work of students attending the meeting, as well as software packages offering robust implementations of Newton–Krylov methods and other nonlinear solvers. The students' topics covered a broad range of areas, including solvers for electrical tunneling, optimization methods for groundwater remediation, globalization of Newton–Krylov methods for the Navier–Stokes equations, and variants of Newton–Krylov methods for problems with expensive nonlinear function evaluations. The software posters provided information on the Sandia package NOX, the Lawrence Livermore package SUNDIALS, and the suite of packages offered through the TOPS (Terascale Optimal PDE Simulations) SciDAC project (www.tops-scidac.org), including the Argonne package PETSc. These packages, all available to the public, offer various functionalities in addition to Newton–Krylov solvers.

Although the presentations at the workshop reflected significant advances in the field of nonlinear solvers, many challenges remain. Newton–Krylov methods have reached an advanced stage of development, as evidenced by their effectiveness across a broad range of applications and by their implementation in sophisticated software packages. Improving the robustness of these methods, however, remains a very active area of research, as do the development and incorporation of new tools and techniques, such as automatic differentiation and multilevel approaches.

Similarly, there have been significant advances in preconditioning, including new application-specific precondi-tioners and new approaches like nonlinear preconditioning; nevertheless, this very important area is likely to see continued development for the foreseeable future. Pseudo-transient and natural-parameter continuation methods were more in evidence at this workshop than at previous ones and seem likely to become more widely applied to large-scale problems. Robust solvers for problems discretized with discontinuous Galerkin methods continue to be an open research area, as does productive use of multiple grid levels in nonlinear solvers. Fault-tolerant implementations of nonlinear solvers also remain as an important open area.

Advances in algorithms for nonlinear solvers are helping to push back the frontiers of science in many areas, and solution methods for large-scale nonlinear problems constitute a very active research area. Interested researchers can learn more about this field in sessions on nonlinear solvers at the Copper Mountain Conference on Iterative Methods, March 28–April 2, 2004, and at a mini-symposium titled "Transitioning Nonlinear, Time-Dependent Codes From Explicit to Implicit Formulations," planned for the 2004 SIAM Annual Meeting in Portland. For more information about the workshop described in this article, see http://www.llnl.gov/casc/workshops/nonlinear_2003, where abstracts and PDF files for many of the presentations are archived.

Acknowledgment

This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract number W–7405–Eng–48. The first author was also supported in part by subcontract B341493, University of California Lawrence Livermore National Laboratory, and by contracts 16099 and 198313 with Sandia National Laboratories, funded in part by the DOE ASCI program and in part by Sandia's Computer Science Research Institute.

Homer Walker, a professor in the Mathematical Sciences Department at Worcester Polytechnic Institute, is currently on a sabbatical visit to Sandia National Laboratories. His main research interests are iterative methods for large-scale linear and nonlinear systems, especially Krylov subspace methods and Newton–Krylov methods. Carol Woodward is a researcher at the LLNL Center for Applied Scientific Computing. Her research interests include nonlinear solvers, especially Newton–Krylov and nonlinear multilevel methods, nonlinear diffusion equations, and parallel computing.