SCALABLE SOLVERS IN DOE'S SciDAC INITIATIVE

David E. Keyes

Department of Applied Physics and Applied Mathematics
Columbia University
New York, NY 10027
kd2112@columbia.edu

The Scientific Discovery through Advanced Computing (SciDAC) initiative of the U.S. Department of Energy is a web of 51 interconnected projects — partly research and partly software development — designed to support simulation, data exploration, and collaboration in many thrust areas including: climate modeling, fusion energy, chemical and materials science, astrophysics, and high energy and particle physics. SciDAC supports the creation of a new generation of scientific simulation codes for terascale systems. The program also includes research on numerical algorithms and systems software that will allow these codes to use modern parallel computers effectively. We very briefly introduce SciDAC and then focus specifically on its Terascale Optimal PDE Simulations (TOPS) project, which engages three DOE laboratories and six universities in the development of scalable software for the integration of PDEs, along with sensitivity analysis, PDE-constrained optimization, and modal or stability analysis of systems governed by PDEs.

SciDAC charts a new direction in federally funded research in several respects. First, it is a bold affirmation of the importance of simulation in leading to new science. In addition to its traditional role of corroborating experiment and theory, simulation can *lead* experiment and theory if a sufficiently highly resolved high fidelity model of the fundamental physics can be constructed and solved. Simulation is an important extension to experimentation, wherever experimentation is controversial, dangerous, illegal, difficult to observe, or simply very expensive! Second, SciDAC recognizes that leading edge computational science is interdisciplinary. It does not provide physics code teams support to write their own solvers or meshers or to develop parallel systems software. These latter capabilities must indeed be *driven* by applications, but developed by specialists, with a premium on reuse across many projects and on extensibility. Third, SciDAC requires joint laboratory and university participation. Universities are strong in generating and testing ideas but poor in developing and maintaining software packages over many generations of students. There is good complementarity in a lab-based, university-shared project. To support such geographically distributed collaboration, SciDAC devotes significant resources to grid software. Fourth, SciDAC commits a generation of community simulation codes to the model of distributed hierarchical memory supercomputers. SciDAC involves 13 DOE laboratories and more than 50 universities throughout the United States, plus a handful of co-PIs abroad.

Thirty-three projects are in the biological, chemical, and physical sciences. Specifically, fourteen projects will advance the science of climate simulation and prediction. These projects involve both novel methods and computationally efficient approaches for simulating components of the climate system and work on an integrated climate model. Ten projects address quantum chemistry and fluid dynamics, which are critical for modeling energy-related chemical transformations such as combustion, catalysis, and photochemical energy conversion. The goal of these projects is efficient computational algorithms to predict complex molecular structures and reaction rates with unprecedented accuracy. Five projects are focused on developing and improving the physics models needed for integrated simulations of plasma systems to advance fusion energy science. These projects will focus on such fundamental phenomena as electromagnetic wave-plasma interactions, plasma turbulence, and macroscopic stability of magnetically confined plasmas. Four projects in high energy and nuclear physics will explore the fundamental processes of nature. The projects include the search for the explosion mechanism of core-collapse supernovae, development of a new generation of accelerator simulation codes, and simulations of quantum chromodynamics (QCD).

Seventeen projects will develop the software infrastructure to support research collaboration using distributed resources and scientific simulation on terascale computers. Four national collaboratory, two middleware, and four network research projects will seek to research, develop, deploy and refine the underpinning software environment that will enable innovative approaches to scientific computing through secure remote access to shared distributed resources, large-scale transfers over high-speed networks, and integration of collaborative tools with the researcher's desktop. Four computer science teams will address critical issues in high performance component software technology, large scale scientific data management, understanding application/architecture relationships for improved sustained performance, and scalable system software tools for improved management and utility of systems with thousands of processors. Finally, three applied mathematics teams will provide interoperable scalable numerical libraries for meshing, discretization, and solution.

Under TOPS, we have been pursuing optimal parallel algorithms for PDE simulations of Newton-Krylov type, preconditioned with Schwarz domain decomposition and multilevel methods. Simultaneous advances in object-oriented software engineering have enabled the construction of complex software systems in which these algorithmic elements can be combined modularly, recursively, and relatively efficiently in parallel. We outline the basic algorithmic methodology for scalable nonlinearly implicit solution of PDEs, including Jacobian-free methods and the roles of automatic differentiation and physics-based preconditioning. We illustrate with recent work in magnetohydrodynamics.

TOPS is unclassified and its codes are freely available, including newly interoperable versions of PETSc, Hypre, SUNDIALS, SuperLU, and other codes. The speaker is the director of the TOPS project and will provide a personal perspective on SciDAC, and connect to minisymposia at USNCCM7 on TOPS, automatic differentiation, and PDE-constrained optimization.

References

[1] D. E. Keyes, "Terascale Implicit Methods for Partial Differential Equations," *Contemporary Mathematics*, v. 306, pp. 29-84, AMS, Providence, 2002.
[2] TOPS, "Terascale Optimal PDE Simulations," project homepage, `www.top-scidac.org`.