CLIMATE MODELING

Climate modeling is perhaps the largest computational challenge the human race has attempted to date. Only a handful of applications exist on a larger scale, and none require the same level of detail. The number of components required to work together is also unprecedented in the field of computational science: atmosphere, land, ocean, and ice models (joined by a flux coupler) must work with radiation, cloud, chemistry, advection, soil, vegetation, and water runoff models (not to mention a whole host of subgrid parameterizations) to produce meaningful results. We can simplify some of these models, depending on the question to be answered, but our climate system's complex interactions will continue to strain the limits of our largest supercomputers for the foreseeable future.

Simulating the Earth's climate

Because we cannot perform large-scale experiments on the Earth's climatological system, these climate models are our only laboratory for getting answers to questions about how our environment behaves and how it responds to changes—human-made or otherwise. The answers we obtain provide the foundation for debate about our ecological policies, which in turn profoundly affect how we live our lives and conduct our business.

Average global temperatures are rising, and our current models support the conjecture that human activity contributes to that rise. Those who oppose policy changes that would address this conclusion point to the inaccuracies in current models to dismiss the results. Implicit in this dismissal is the assumption that the errors bias results toward a warmer climate. However, even if we choose to ignore collective scientific and model predictions, we cannot ignore the warming trend and the commonsense conclusion that humans are likely contributors. In any event, between humans and nature, only humans can intentionally attempt to reduce the warming trend. To that end, climate models provide a relatively new and extremely valuable tool.

The development of climate models, especially the atmospheric component, is somewhat unique in the field of numerical modeling. Most computational fluid dynamicists are surprised to learn, for example, that almost all production-level atmospheric models are based on a spectral expansion, not the more common finite volumes or finite elements. Spectral approximations are rarely even considered for most other applications, because they are simply unworkable for most geometries. The atmosphere's geometry, however, is a simple spherical shell, its natural coordinate system is spherical, and the spherical coordinate system has built-in singularities at the poles. For example, the total derivative of the velocity is infinite at the poles. No approach handles these singularities as elegantly as the spherical harmonic spectral transform method. Also, when spectral models came to dominate, the most powerful computers consisted of vector processors—and the spectral method vectorizes beautifully.

Computer chip design is dictated by market

1521-9615/02/$17.00

WILLIAM F. SPOTZ
Sandia National Laboratories
PAUL N. SWARZTRAUBER
National Center for Atmospheric Research

forces, which care nothing for what will most efficiently run our most elegant atmospheric models. So, climate modelers and weather forecasters must use distributed-memory architectures that are ill-suited for the types of models available to them. Indeed, although processing loads for spectral models can be balanced for specific combinations of resolution and numbers of processors, the overhead due to data transpositions causes a precipitous drop-off in scaling at what would be considered a low to moderate number of processors by today's standards.

Thus, to best use the supercomputers available to us, we must change current numerical methods. This change (which is admittedly meeting some resistance within the atmospheric community), coupled with increasing computational power, should one day provide capabilities that would never be possible with the spectral method. As an example, *local methods* provide the capability to refine the solution around discontinuities in the problem such as those introduced by the Earth's topography, weather fronts, or clouds.

In this issue

This special issue on climate modeling presents a snapshot of some of the research scientists are conducting to improve our climate-modeling capabilities. We present articles on promising atmospheric and ocean models, meshing issues for geophysical applications, implementation on distributed computer architectures, atmospheric-chemistry models, and the testing and evaluation of models. These articles, taken together, provide a good indication of the problems facing climate modelers today and some of the solutions they are pursuing.

Despite the breadth of the articles in this special issue, we were unsuccessful in obtaining invited articles about the spectral method, mesh refinement, ensemble modeling, subgrid parameterizations, turbulence, convergence, ice models, or flux couplers, nor a comprehensive overview of climate supercomputing centers. Had we succeeded in getting articles on some or most of these topics, the magazine would not have had sufficient space to accommodate them. A truly comprehensive overview of climate modeling would require more than a single issue of *CiSE*. So, what you have before you is an intentionally incomplete collection of articles, slightly biased toward our own personal research area (the approximation of the solution to partial differential equations in spherical geometry). Nevertheless, we believe this issue will inform the interested technical reader about how some things are done and where things are headed in the field of climate modeling.

Bill Spotz is a senior member of the technical staff at Sandia National Laboratories in Albuquerque, New Mexico. His research interests include efficiency and accuracy issues related to nonlinear solvers, and discretization methods for the solution of partial differential equations. He has a PhD in aerospace engineering from the University of Texas. He is a member of AIAA and SIAM. Contact him at Sandia National Laboratories, PO Box 5800, Albuquerque, NM 87185-1110; wfspotz@sandia.gov.

Paul Swarztrauber is a senior scientist in the Scientific Computing Division at the National Center for Atmospheric Research and professor adjoint in the Computer Science Department at the University of Colorado. His research interests include elliptic and hyperbolic partial differential equations, computational fluid dynamics, linear algebra, parallel computation and communication algorithms, computational aspects of weather and climate modeling, vector and scalar harmonic transforms, the fast Fourier transform, and computational methods for vector and multiprocessor computers. He has a BS in engineering physics from the University of Illinois and an MS in applied mathematics and a PhD in mathematics, both from the University of Colorado. Contact him at the Nat'l Center for Atmospheric Research, PO Box 3000, Boulder, CO 80307-3000; pauls@ucar.edu.