Discovery at the Interface of Science and Engineering: **Science Matters!**

Computer and Information Sciences
Modeling and Simulation

Community Climate System Model Development for Petascale Supercomputers

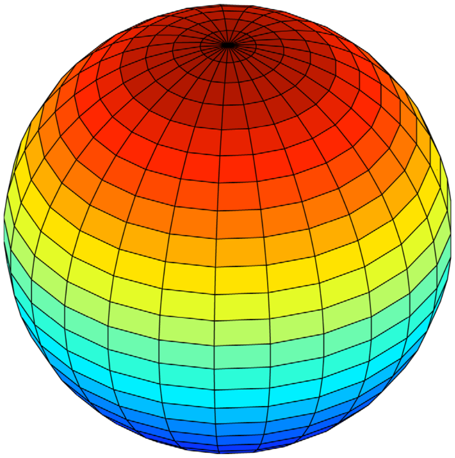

Figure 1: A latitude-longitude grid showing the clustering of grid points at the poles. This creates numerical difficulties in climate models which can be overcome, but at the cost of degrading parallel scalability.

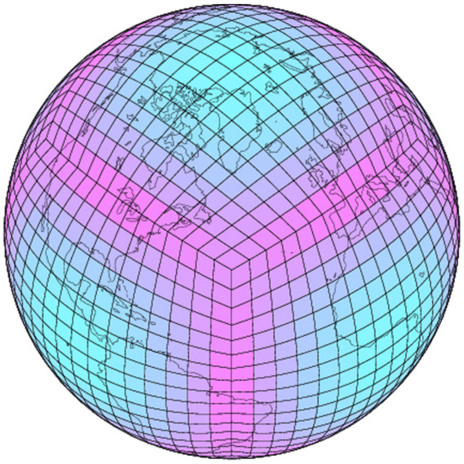

Figure 2: The cubed-sphere grid used by the spectral element atmospheric model component of the CCSM. The color represents the area of each grid cell. The cells are much more uniform when compared to the cells in a latitude-longitude grid.

Predicting local impacts of climate change with next-generation computers

For more information:

Technical Contact:
Mark Taylor
505-284-1874
mataylo@sandia.gov

Science Matters Contact:
Alan Burns, Ph.D
505-844-9642
aburns@sandia.gov

Climate change is already underway and will accelerate this century, but there are still major uncertainties. Human societies will respond with both mitigation and adaptation strategies, which may include changes in consumption of energy and other resources, land use, agriculture, and migration. These stresses could also lead to a shift in alliances and to civil unrest and conflict. Adapting to new situations will require advanced understanding of climate change and its direct and indirect effects. An improved understanding of climate change is of paramount importance for mission planning, surveillance, facilities placement and design, and sustainable energy futures. Developing this understanding requires high fidelity climate models which simulate many components of the Earth system, including atmospheric and oceanic circulation, land surface processes, sea ice, and chemical and biogeochemical cycles. Adding full chemical and biogeochemical cycles to today's climate models will require significantly more computational resources than available today. Furthermore, an increase in resolution to allow for the accurate assessment of regional impacts of climate change also requires further increases in computational power. Achieving both of these goals will require petascale computer architectures.

The Department of Energy's (DOE) upcoming petascale computers will have hundreds of thousands of processors. Effectively using such machines remains a challenge due to several scalability bottlenecks present in all modern climate models, the largest of which is created by the numerical methods used in the *dynamical core* of the atmospheric model component. The dynamical core solves the partial differential equations governing the fluid dynamical aspects of the atmosphere. In addition, atmospheric models contain a suite of subgrid parametrizations for the many physical processes unresolved in an atmospheric model but which drive the dynamics such as convection, precipitation

Sandia National Laboratories

National Nuclear Security Administration

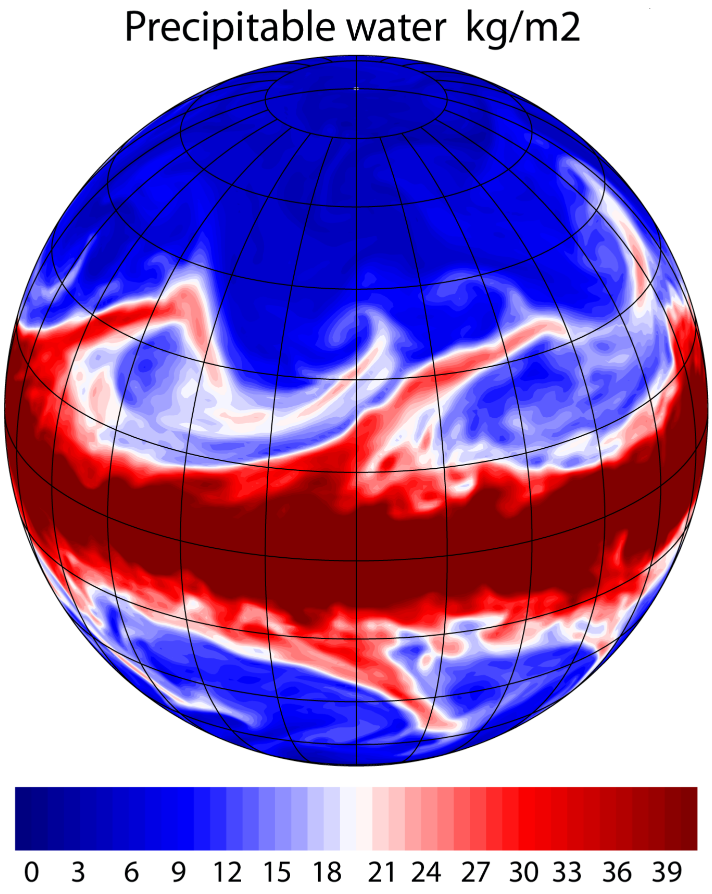

Figure 3: A snapshot of precipitable water over the surface of an Aqua planet, simulated using the atmospheric component of the CCSM with the spectral element dynamical core on a cubed-sphere grid. Aqua planet simulations are used to test new dynamical cores and other physical processes in climate models before coupling with other components such as land, ocean and ice models. The color scheme is suggestive of clouds, but we note that even on petascale computers, climate models will lack sufficient resolution to directly model clouds. Instead, the effects of clouds on the resolved scales are modeled with subgrid parametrizations.

and radiative forcings. Currently, most dynamical cores use latitude-longitude based grids (Fig. 1). These grids create a logically Cartesian orthogonal mesh suitable for a wide array of numerical methods, including finite volumes and the spectral transform method. The grid lines cluster at the pole, creating several computational difficulties collectively referred to as the *pole problem*. One such difficulty comes from the Courant-Friedrichs-Lewy restriction, which tells us that this clustering requires the dynamical core to compute many more timesteps than would otherwise be necessary. This would result in a simulation that should take days to complete now requiring a few months instead. There are many successful techniques to handle this pole problem, however most of them substantially degrade parallel scalability by requiring too much inter-processor communication.

Team members at Sandia, the National Center for Atmospheric Research, and Oak Ridge National Laboratory, all lead by Mark Taylor, have thus been focusing on the development of new, more scalable dynamical cores based on cubed-sphere grids (Fig. 2) for the Community Climate System Model (CCSM), the U.S. flagship climate change model. They have recently developed a new formulation of the highly scalable spectral element dynamical core tailored to the needs of the CCSM. This work has lead to unprecedented scalability in the atmospheric component. It can now run efficiently on 86,000 processors when using a horizontal average grid spacing of 25 km. Even better scalability will be possible when computing with a global resolution of 10 km, DOE's long term goal. The team has completed extensive verification work using standardized atmospheric tests with prescribed surface temperatures but without the CCSM land, ice or ocean models (Fig 3). The team is currently focused on coupling with these other CCSM component models.

Sandia National Laboratories

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND2008-5080P

NNSA National Nuclear Security Administration

08/2008