Global Weather Prediction and High-End Computing at NASA

The authors demonstrate the current capabilities of NASA's finite-volume General Circulation Model in high-resolution global weather prediction and discuss its development path in the foreseeable future. This model is a prototype of a future NASA Earth-modeling system intended to unify development activities across various disciplines within NASA's Earth Science Enterprise.

NASA's goal for an Earth-modeling system is to unify the model development activities that cut across various disciplines in the Earth Science Enterprise, which is a NASA organization for all NASA's activities related to Earth science. Earth-modeling system applications include, but are not limited to, weather and chemistry climate-change predictions and atmospheric and oceanic data assimilation. High-resolution global weather prediction, among these applications, requires the highest temporal and spatial resolution, and, hence, demands the most capability from a high-end computing system.

In the continuing quest to improve and perhaps push the limit of weather prediction (see the "Weather Predictability" sidebar), we are adopting more physically-based algorithms with much higher resolution than that of earlier models, which is crucial to improving forecast skill. We also are including additional physical and chemical components such as a chemical transport model previously not coupled to modeling systems.

Because a comprehensive high-resolution Earth-modeling system requires enormous computing power, we must design all component models efficiently for parallel computers with distributed-memory platforms. To this end, we started developing the finite-volume General Circulation Model (fvGCM) of the atmosphere, which is based on the work of Shian-Jiann Lin and Richard Rood[1–4] and their collaboration with the National Center for Atmospheric Research (NCAR). Some of the fvGCM's more technical aspects and climate characteristics appear elsewhere.[5]

In this article, we will first demonstrate the model's current capabilities in high-resolution global weather forecasting by predicting real weather events in terms of both accuracy and efficiency, and then outline the model's development-evolution path and its computer requirements in the foreseeable future.

1521-9615/04/$20.00 © 2004 IEEE
Copublished by the IEEE CS and the AIP

Shian-Jiann Lin, Robert Atlas, and Kao-San Yeh
NASA Goddard Space Flight Center

The Current High-End Modeling System for Weather Prediction

The fvGCM features a unique finite-volume dynamics system with local conservation and monotonicity to ensure a global consistency of simulated or predicted atmospheric dynamical processes. It describes the Earth's surface with the traditional latitude–longitude grid system consisting of a set of grid boxes defined along the latitude circles and along the meridians. We assume that the model atmosphere is in hydrostatic equilibrium—that is, the

Weather Predictability

Mathematically speaking, numerical weather prediction is an initial-value problem. Assuming all equations governing the atmosphere's motions are known and can be solved exactly—for example, assuming we have a perfect model—the prediction of a future state of the atmosphere relies entirely on the correctness of the input initial conditions. Because the initial conditions for the atmosphere and the ocean will never be perfectly prescribed, the weather's predictability is limited by a certain amount of error in the initial conditions, as well as by the chaotic nature of dynamics and physics, which further amplifies the initial errors.

However, even today's most advanced data-assimilation systems still possess significant errors in providing initial conditions, due to imperfections in observation facility and data-assimilation techniques. Furthermore, our understanding of the dynamical, physical, chemical, and biological processes of the Earth environment still is far from complete, and our modeling techniques are still far from perfect. Thus, we are not yet even close to the theoretical limit of weather predictability; we have much still to gain by improving our systems. For more information about atmospheric modeling, data assimilation, and predictability, please read Eugenia Kalnay's textbook on the subject.[1] (Kalnay served for many years as the director of the Environmental Modeling Center [EMC]. The National Weather Service's weather forecasts are based on EMC's model predictions.)

Two major areas need improvement: numerical approximations to the dynamical and the physical processes, respectively. We've known the analytic equations governing the fluid-dynamical processes of the atmosphere and the ocean for more than a century; it is the numerical solutions to these well-known and trusted equations that advanced numerical algorithms and increasing resolution can improve. The errors in the parameterized physical processes, however, cannot be reduced by simply increasing the resolution, because some of the physical processes, such as moist convection for the formation of clouds and the associated cloud-radiation processes, are either not well understood or not fully described by existing equations. In particular, cumulus scales are still not predictable beyond a few hours, leading to the need for probabilistic forecasting. Increasing the resolution, however, can reduce the reliance on physical parameterizations and lead to the direct use of explicit formulation for crucial physical processes. Therefore, our approach in the current and future modeling system is to increase resolution to the maximum extent allowed by available computer platforms and seek a direct, physically-based approach to modeling the physical processes at that resolution.

Reference

1. E. Kalnay, *Atmospheric Modeling, Data Assimilation and Predictability*, Cambridge Univ. Press, 2003.

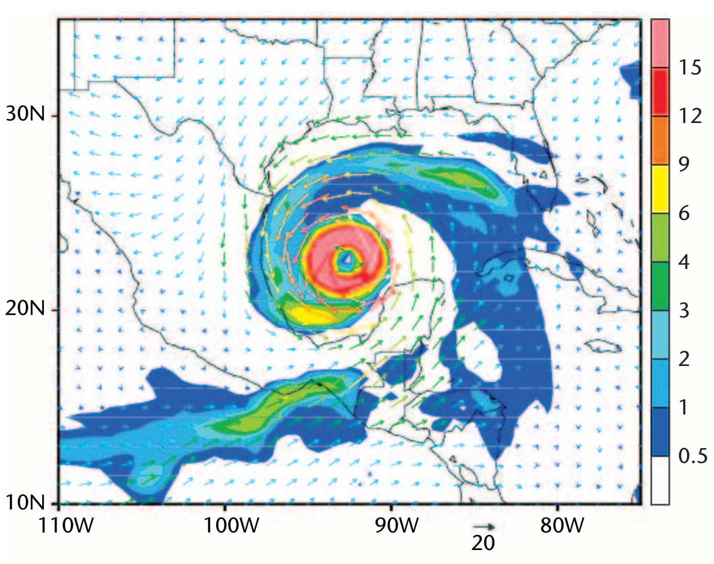

Figure 1. A hurricane simulated by the NASA finite-volume General Circulation Model. The precipitation rate (mm/hr) is depicted with the color scheme on the right, and the wind (m/s) at the 850 millibar (mb) level is shown with magnitude proportional to the arrow (20 m/s) at the bottom.

air's gravity is balanced by the air pressure. We decouple the model atmosphere vertically into a sequence of horizontal layers that are described vertically with Lagrangian coordinates—that is, the altitudes of the layers that evolve with time.[4] The model atmosphere evolves by integrating the conservation laws of physics over the finite volumes specified by the grid boxes within each layer. The total air mass and the total energy are thus conserved in the dynamics system. (You can find more details of the fvGCM's design at our Web site: http://polar.gsfc.nasa.gov/sci_research/atbd_pages/nextgeneration.php.)

The fvGCM

We produced all model simulations in this article with a resolution of 0.625 degrees longitude, 0.5 degrees latitude, and 32 layers in the vertical, covering the atmosphere from the ground surface to an altitude of approximately 55 km.

Figure 1 depicts the precipitation and low-level wind structures of a hurricane produced by the model, over the Gulf of Mexico, in a climate simulation during the early stages of the fvGCM development. This is a fairly large hurricane with the eye clearly defined by the vortex structure and the precipitation pattern with well-pronounced spiral bands. Encouraged by the model's capability to simulate realistic hurricanes and other fine-scale weather events, such as surface fronts and severe winter storms, we proceeded to develop a next-generation data-assimilation system that modifies model predictions with observation data, the finite-volume Data Assimilation System (fvDAS), which is based on the fvGCM and is now operational.

Figure 2 shows how predictions from the fvGCM and the fvDAS were proven correct with the track of Hurricane Floyd, which touched the US mainland in September 1999. The fvDAS initialized the fvGCM at 00 GMT 12 September for a five-day forecast at the resolution of 0.625 by 0.5 degrees. Both the fvGCM forecast (purple squares) and the fvDAS simulation (blue crosses) match very well with the observed best track (red spiral spots) from the US National Hurricane Center.

In spring 2002, motivated by the initial success of the fvGCM's forecasts of hurricanes (Floyd and many others), we began evaluating the fvGCM's capability with 10-day global weather forecasts in real time. We found the model's skill at forecasting to be very competitive with many US operational centers. Figures 3 and 4 show a few recent severe storms. They depict, at time of forecast, the sea-level pressure, the instant precipitation rate, and the total snow accumulation on the ground. Figure 4 also includes the winds 10 meters above the ground.

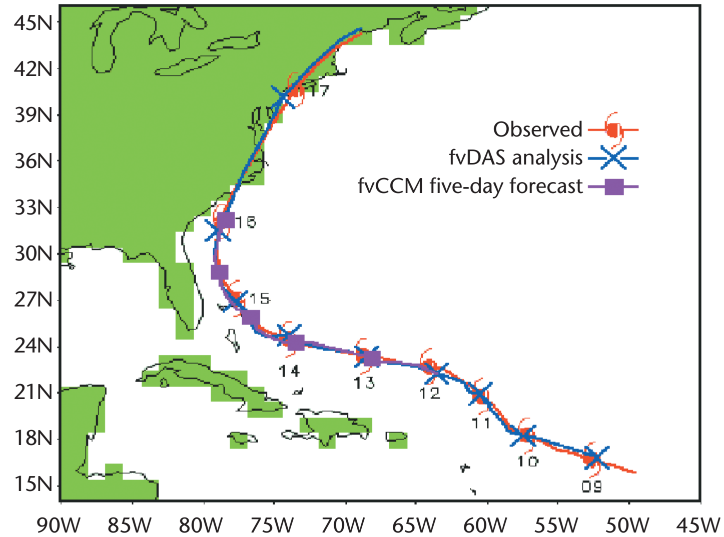

Figure 2. Validation of the NASA finite-volume General Circulation Model (purple squares) and the data-assimilation system (blue crosses) with the track of Hurricane Floyd (red spiral spots) observed by the National Hurricane Center in September 1999.

Figure 3 demonstrates the fvGCM's accuracy in forecasting severe snowstorms over the Mid-Atlantic states, such as Maryland and Virginia. Figure 3a shows the model's initial state at 03 GMT on 13 February. Figure 3b shows the model prediction after 5.5 days, revealing a 24-inch snow accumulation for the Baltimore–Washington area

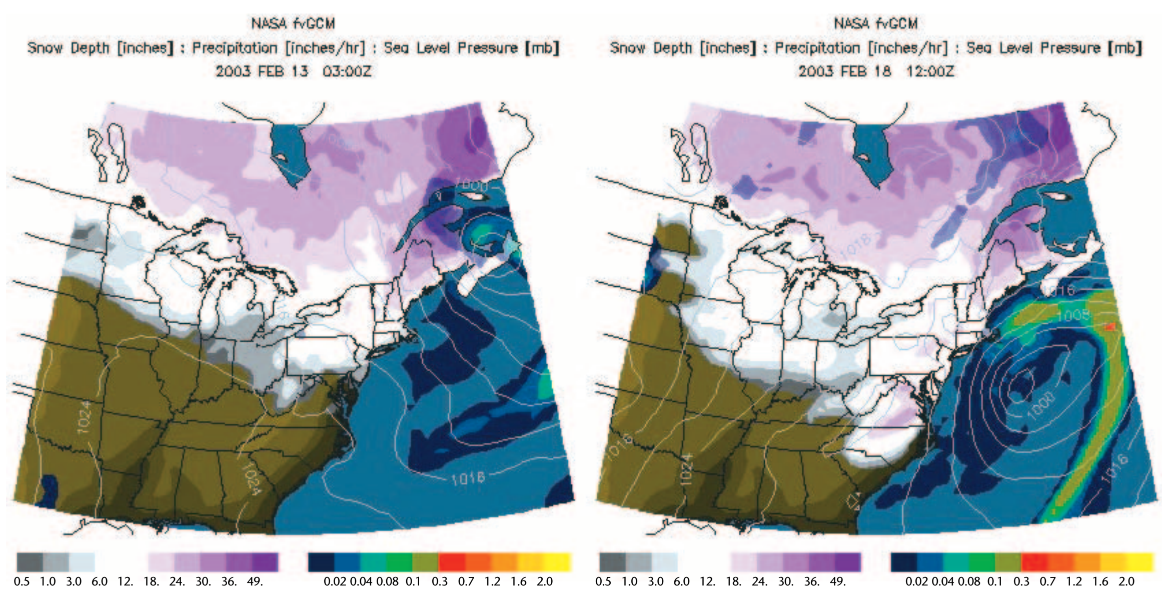

Figure 3. Snow accumulation. (a) The initial state at 03 Greenwich Mean Time (GMT) 13 February 2003 (sea-level pressure, instant precipitation rate, and the total snow accumulation on the ground) for a snowstorm on the US East Coast in February 2003. (b) As in (a), but for the 5.5-day forecast at 12 GMT 18 February 2003.

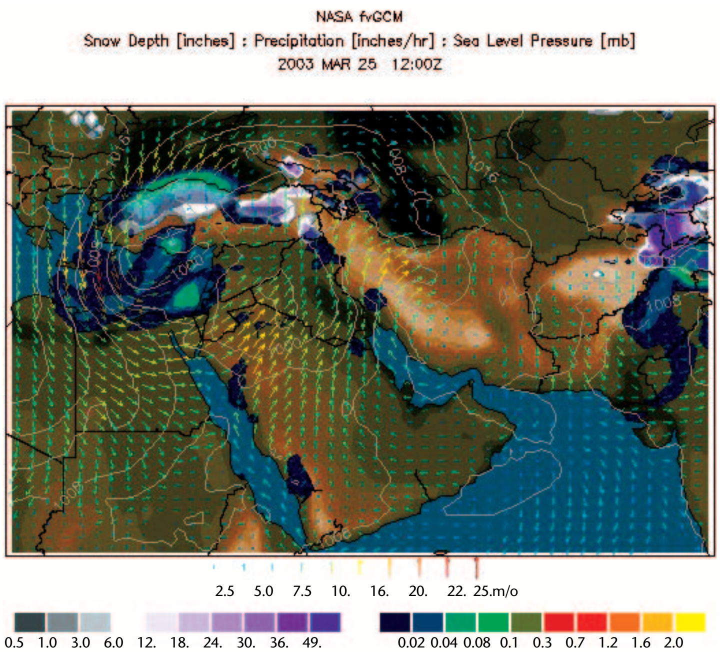

Figure 4. A 5.5-day (00 GMT 20 March to 1200 GMT 25 March) forecast of a severe snow- and sandstorm in the Middle East during the US's Operation Iraqi Freedom in 2003.

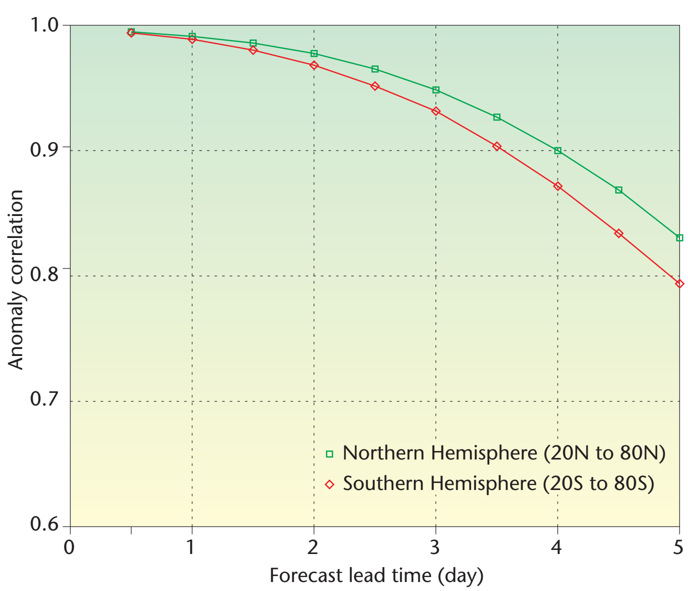

Figure 5. The 500-millibar (mb) height anomaly correlation between the NASA finite-volume General Circulation Model and an analysis by the National Centers for Environmental Prediction from 1 December 2002 to 28 February 2003. The anomaly correlation in the Northern Hemisphere (winter) is depicted with green rectangles; the Southern Hemisphere (summer) with red diamonds.

with amounts up to and beyond 30 inches in the mountains of western Maryland and Virginia. The snow accumulation over the entire US East Coast is also quite accurate. In fact, we forecasted this particular "storm of the new century" for the Mid-Atlantic states 10 days in advance (in private email exchanges between NASA and National Oceanic and Atmospheric Administration [NOAA] scientists). However, because we are not to interfere with the National Weather Service, this prediction went unannounced.

Accurate and timely weather forecasts can be useful for military strategies. Immediately after US Operation Iraqi Freedom in March 2003, we added a window to highlight our real-time global weather forecast in the Middle East. We initiated this model forecast at 00 GMT on 20 March, and the model successfully predicted a severe storm in Turkey and the Black Sea, as we describe later, snow in northern Turkey, and a sandstorm over the Arabian desert that occurred a week later. Figure 4 shows a mature severe storm on the sixth day of the Iraq war that the fvGCM forecasted. In the figure, the storm approaches from the Mediterranean and Black Sea with strong winds up to 15 meters per second (m/s; 34 miles per hour) in the southwest deserts of Iraq and northern Saudi Arabia. The wind continues to intensify to up to 23 m/s (52 miles per hour) the next day as the storm moves east, causing brutal sandstorms in the deserts for 48 hours—exactly as described later by a CNN live report. For more details of the NASA fvGCM real-time global weather forecast, visit our Web site at http://polar.gsfc.nasa.gov/sci_research/fvdas/NASCAR_web/nwp.

To illustrate the fvGCM's general accuracy in global weather forecast, Figure 5 summarizes its forecast skills with a 90-day average of anomaly correlation—the correlation between model forecast and observation relative to climatology. This anomaly correlation is averaged over the period from 1 December 2002 to 28 February 2003, and is calculated at the 500-millibar (mb) pressure level, which represents the middle troposphere. According to a US National Centers for Environmental Prediction analysis, the fvGCM shows very good forecast skills for up to five days in forecast time, with a score of 83 percent in the Northern Hemisphere (winter) and 79 percent in the Southern Hemisphere (summer) for the fifth forecast day.

Real-time weather forecasts are time-critical, and horizontal resolution, or the horizontal distance between data points, is crucial to global models' capability to predict severe weather events. However, the time it takes to finish a forecast increases at least quadratically with each doubling of the horizontal resolution, and need additional overhead when using smaller time steps to stabilize the horizontal dynamics of the fvGCM when increasing the horizontal resolution. Therefore, NASA has invested substantial resources in the software engineering and optimization of the fvGCM modeling system. Figure 6 shows the fvGCM's computational performance on an IBM RS/6000 SP system named Eagle (purple squares), on an SGI Origin 3000 system named Daley (blue diamonds), and on a Compaq AlphaServer SC45 system named Halem (red triangles). The abscissas refer to the number of CPUs employed during the computation, and the ordinates correspond to the throughput—namely, the number of simulation days carried out by the model per wall-clock day (real time). The increasing rate of throughput with respect to the increasing number of CPUs is called the *parallel efficiency* of the model for the given resolution. The model's throughputs on all three machines increase rapidly with the increasing number of CPUs up to about 250, indicating a quite efficient parallel implementation at the operational resolution of 0.625 by 0.5 degrees. Although the SGI machine appears to have a slightly better scaling of linear throughput, the processor speed is about 2.5 times faster on Halem, hence it is not surprising to see that Halem outperforms Daley in the low- to mid-range CPU counts (32 to 256).

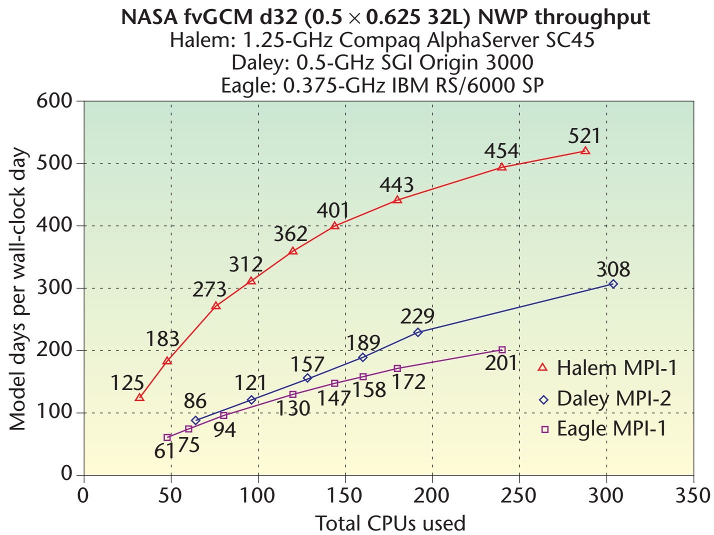

Figure 6. Computational performance of the NASA finite-volume General Circulation Model on various machines, including IBM (purple squares), SGI (blue triangles), and Compaq (red triangles).

Future Development

The improvement in parallel efficiency and advancement in scientific algorithms can bring as much improvement to the overall computational efficiency as the hardware improvement predicted by Moore's law (which states that computing power doubles about every 18 months). Although the current fvGCM's parallel efficiency is adequate for today's high-end computers, it won't remain so. As Figure 6 shows, the fvGCM's parallel efficiency (slope) on Halem for the current operational resolution (0.625 by 0.5 degrees) reduces to that of a lower-frequency Daley at about 250 CPUs. Today's high-end computers can have as many as 4,000 CPUs in a system (see, for example, the Japanese Earth Simulator). To meet, or even beat, the Earth Simulator's computing power, we might expect as many as 50,000 CPUs in a high-end parallel computing system in the US by 2010. This poses a serious challenge to our ability to improve the model's parallel efficiency.

The current model's limits in parallel efficiency are mainly due to the use of the traditional latitude–longitude grid system in which the meridians converge at the computational poles. This convergence makes a 2D domain decomposition technique less efficient, even undesirable. To resolve the issue of degraded parallel efficiency caused by the presence of computational poles, it seems wisest to abandon the traditional latitude–longitude grid system and seek a quasi-uniform grid system that has no computational poles. For quasi-uniform grid systems, however, formulating accurate high-order numerical solutions to the physical laws that govern the atmosphere is very difficult. Because of the lack of orthogonality within the grid system, it is even more difficult to maintain the fvGCM's dynamics features of local conservation and monotonicity.

Ross Heikes and David Randall[6] investigated the convergence problem (which is when the error of a numerical solution approaches to zero as the distance between data points goes to zero) of numerical solutions on an icosahedral geodesic grid. Hirofumi Tomita and his colleagues[7] modified the standard icosahedral grid with spring dynamics and gravitational centers to minimize grid-related truncation errors of numerical approximations. We have adopted the spring-dynamics geodesic grid Tomita designed for a possible future generation of the fvGCM (see Figure 7). We also are looking into the possibility of using a cubed sphere, which is obtained by projecting a cube from Earth's center onto the spherical surface. This helps avoid the computational poles mentioned earlier. Icosahedral geodesic grids and cubed spheres are the two most popular approaches for that purpose. To maintain the current fvGCM's exceptional quality in a future modeling system, we are designing new computational algorithms with the desired properties of local conservation and monotonicity for quasi-uniform grid

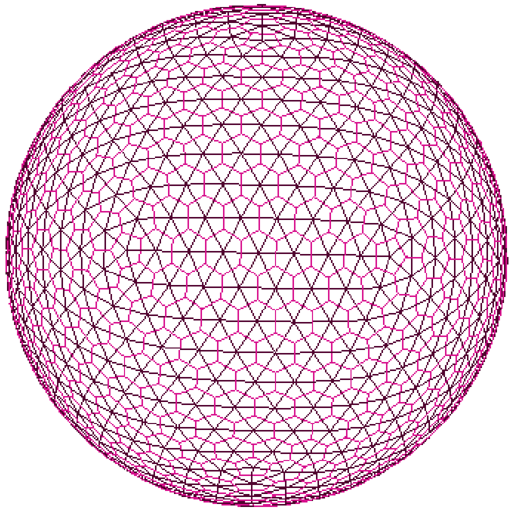

Figure 7. A spring-dynamics icosahedral geodesic grid with gravitational centers for global high-resolution modeling. Parallel computation is efficient on this grid, because there are no computational poles. Errors of numerical approximations are reduced by the smoothness of the grid lines (black lines) and the definition of the data points (intersections of black lines) at the gravitational centers of the grid cells (magenta lines).

systems. We've built a prototype of our new finite-volume formulation with the desired numerical properties, and it seems that we will be able to thoroughly resolve any parallel efficiency issues associated with the fvGCM's current design. The goal is to fully utilize available computing power for atmospheric modeling with global high resolution of a few kilometers between 2010 and 2015.

An atmosphere model alone is not sufficient to make realistic weather and climate predictions. We need land, ocean, and sea-ice models to provide conditioning (in the form of momentum, latent heat, and sensible heat fluxes) to the lower atmosphere model's boundary. An ocean model needs fluxes from an atmosphere model to drive and, therefore, predict ocean currents. A unified modeling system with high parallel-computing efficiency is thus desirable to address the prediction and environmental issues on Earth as well as on other planets, such as Mars. For these purposes, we are proposing an ambitious scientific and engineering project to construct a comprehensive modeling system—Virtual Planet—and to build an exceptionally powerful supercomputer, Planet Simulator.

Most of the uncertainty in predicting the weather and climate stems from the inadequacy of cumulus parameterization—the modeling treatment for the effects of cumulus clouds that are not resolvable with today's computing power. Virtual Planet's ultimate goal is to explicitly formulate the cloud processes, to avoid the uncertainty due to cumulus parameterization. To this end, we'll need a horizontal resolution of 5 km or finer, which could take as many as 50,000 of the most advanced US-made microprocessors to build the massively parallel Planet Simulator. To achieve the ambitious goal of ultra-high global resolution with Planet Simulator, we will need a corresponding upgrade of all components for Virtual Planet. This implies a tremendous amount of research with high-level difficulties, such as

- development of nonhydrostatic finite-volume dynamics of high-order accuracy,
- development of cloud microphysics without cumulus parameterization,
- development and coupling of an eddy-resolving ocean model,
- development and coupling of a dynamical sea-ice model,
- development and coupling of an ultra-high-resolution land model,
- development and coupling of a full atmospheric chemistry model, and
- assimilation of NASA and NOAA high-resolution satellite data.

A project this scale will likely require a coordinated national effort involving several agencies, such as the US Department of Energy, NASA, NOAA, and the National Science Foundation, and research institutions and universities across the US.

Acknowledgments
We thank our colleagues, William Putman, Jiun-Dar Chern, Bo-Wen Shen, and Joseph Ardizzone, for providing useful information and preparing the figures presented in this article. We also appreciate the comments and suggestions from the editors and several anonymous reviewers. This work is sponsored by the NASA Earth Science Enterprise Science Division through the Global Modeling and Analysis Program.

References
1. S.-J. Lin and R.B. Rood, "Multidimensional Flux-Form Semi-Lagrangian Transport Schemes," *Monthly Weather Rev.*, vol. 124, no. 9, 1996, pp. 2046–2070.
2. S.-J. Lin and R.B. Rood, "An Explicit Flux-Form Semi-Lagrangian Shallow-Water Model on the Sphere," *Quarterly J. Royal Meteorological Soc.*, vol. 123, no. 544, 1997, pp. 2477–2498.

3. S.-J. Lin, "A Finite-Volume Integration Method for Computing Pressure Gradient Force in General Vertical Coordinates," *Quarterly J. Royal Meteorological Soc.*, vol. 123, no. 542, 1997, pp. 1749–1762.
4. S.-J. Lin and R.B. Rood, "Development of the Joint NASA/NCAR General Circulation Model," *Proc. 13th Conf. Numerical Weather Prediction*, Am. Meteorological Soc., 1999, pp. 115–119.
5. K.-S. Yeh, S.-J. Lin, and R.B. Rood, "Applying Local Discretization Methods in the NASA Finite-Volume General Circulation Model," *Computing in Science & Eng.*, vol. 4, no. 5, 2002, pp. 49–54.
6. R.H. Heikes and D.A. Randall, "Numerical Integration of the Shallow-Water Equations on a Twisted Icosahedral Grid: I. Basic Design and Results of Tests," *Monthly Weather Rev.*, vol. 123, no. 6, 1995, pp. 1862–1880.
7. H. Tomita et al., "Shallow Water Model on a Modified Icosahedral Geodesic Grid by Using Spring Dynamics," *J. Computational Physics*, vol. 174, no. 2, 2001, pp. 579–613.

Shian-Jiann Lin is a meteorologist at the NASA Goddard Space Flight Center. His research interests include numerical algorithms for geophysical fluid dynamics and cloud physics. He received a BS in agricultural mechanical engineering from National Taiwan University, an MS in aerospace and mechanical engineering from the University of Oklahoma, and an MA in geophysical fluid dynamics and a PhD in atmospheric science from Princeton University. Contact him at Code 900.3, NASA Goddard Space Flight Center, Greenbelt, MD 20771; lin@gmao.gsfc.nasa.gov.

Robert Atlas is the chief meteorologist at the NASA Goddard Space Flight Center and an adjunct professor at the University of Maryland. He has performed research for applying space-based observations to atmospheric science since 1976, and was the first person to demonstrate the beneficial impact of satellite temperature soundings and surface wind data on weather prediction. He received a BS in aeronautics from St. Louis University, and an MS in meteorology and a PhD in meteorology and oceanography from New York University. He is a member of the Council of the American Meteorological Society and the Scientific Steering Committee for the US Weather Research Program. Contact him at Code 910, NASA Goddard Space Flight Center, Greenbelt, MD 20771; atlas@gmao.gsfc.nasa.gov.

Kao-San Yeh is an assistant research scientist at the University of Maryland and the NASA Goddard Space Flight Center. His research interests include numerical methods for geophysical fluid-dynamics modeling, especially the finite-volume method for conservative and monotonic transport in atmospheric models. He received a BS in physics from National Taiwan University, and an MS in mathematics and a PhD in atmospheric science from Purdue University. Contact him at Code 900.3, NASA Goddard Space Flight Center, Greenbelt, MD 20771; kYeh@gmao.gsfc.nasa.gov.