Global Weather Prediction and High-End Computing at NASA

by

Shian-Jiann Lin, Robert Atlas, *and* **Kao-San Yeh***

NASA Goddard Space Flight Center

**Corresponding author address:* Dr. Kao-San Yeh
Code 900.3, NASA Goddard Space Flight Center, Greenbelt, MD 20771
E-mail: kyeh@gmao.gsfc.nasa.gov

August 18th, 2003

Abstract

We demonstrate current capabilities of the NASA finite-volume General Circulation Model in high-resolution global weather prediction, and discuss its development path in the foreseeable future. This model can be regarded as a prototype of a future NASA Earth modeling system intended to unify development activities cutting across various disciplines within the NASA Earth Science Enterprise.

1. Introduction

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

In the continuing quest to improve and perhaps push to the limit of the predictability of the weather (see the related side bar), we are adopting more physically based algorithms with much higher resolution than those in earlier models. We are also including additional physical and chemical components that have not been coupled to the modeling system previously. As a comprehensive high-resolution Earth modeling system will require enormous computing power, it is important to design all component models efficiently for modern parallel computers with distributed-memory platforms. To this end, we have started with developing the finite-volume General Circulation Model (fvGCM) of the atmosphere, which is based on the work of Lin and Rood [1-4] in 1990s and the collaboration with the National Center for Atmospheric Research (NCAR). Some technical aspects and climate characteristics of the NASA fvGCM are illustrated by Yeh et al. [5]. In this article, we will first demonstrate the model's current capabilities in high-resolution global weather forecast with real weather events in terms of both accuracy and efficiency, then outline the model development-evolution path and the computer requirements in the foreseeable future.

Weather Predictability

Mathematically speaking, numerical weather prediction can be considered as an *initial-value problem.* Assuming all equations governing the motions of the atmosphere are known and can be solved exactly, i.e., assuming a perfect model is given, the prediction of a future state of the atmosphere would then rely entirely on the correctness of the input initial conditions. Since the initial conditions for the atmosphere and the ocean could likely never be perfectly prescribed, the predictability of the weather would be limited by certain amount of error in the initial conditions, as well as by the chaotic nature of dynamics and physics that amplifies the initial errors. However, it is believed that even today's most advanced data assimilation systems still possess significant errors in providing the initial conditions, due to imperfectness of observation facility and data assimilation techniques. Furthermore, our understanding of the dynamical, physical, chemical and biological processes of the Earth environment is still far from being complete, and our modeling techniques are still far from being perfect. Thus we are not yet even close to the theoretical limit of weather predictability, and much can still be gained by improving the modeling/predicting system and the data assimilation system that provide the best estimate of the state of the atmosphere. We describe in this article our continuing efforts towards the improvement of the modeling/predicting system. For more information about atmospheric modeling, data assimilation and predictability, the readers are referred to the textbook by Prof. Eugenia Kalnay [6].

There are two major areas where improvements can be made, in the numerical approximations to the dynamical and the physical processes, respectively. The analytic equations governing the fluid-dynamical processes of the atmosphere and the ocean have been known for more than a century, it is the numerical solutions to these well-known and well-trusted equations that can be improved by advanced numerical algorithms and by increasing the resolution. The errors in the parameterized physical processes, however, can not 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 some crucial physical processes, e.g., cloud microphysics instead of cumulus parameterization, although simulation of precipitation will still require parameterization of the microphysical processes that govern the evolution of the droplet size spectrum. Therefore, our approach in the current and the future modeling system is to increase the resolution to the maximum extent allowed by available computer platforms, and seek for a direct, physically based approach to modeling the physical processes at that resolution.

2. The current high-end NASA modeling system for weather predictions

The NASA fvGCM features a unique finite-volume dynamical system with the properties of local conservation and monotonicity to ensure global consistency of simulated or predicted atmospheric dynamical processes. The model atmosphere is vertically decoupled into a sequence of horizontal layers, within which the latitude-longitude grid system is adopted to approximate 2D conservation laws of physics. Hydrostatic equilibrium is hypothesized within each layer to balance the gravity and the pressure, and the layers are described vertically with Lagrangian coordinates, i.e., the altitudes of the layers which evolve with time [4]. The air mass and total energy are conserved in the 3D discrete dynamical system. For more details of the fvGCM's design, the readers are referred to the documentation provided at our web site <http://polar.gsfc.nasa.gov/sci_research/atbd_pages/nextgeneration.php>.

The fvGCM was initially developed with physical parameterizations from the NCAR Community Climate Model. It was therefore unclear if the model would perform well for high-resolution weather applications. We demonstrate, by examples, the fvGCM's capabilities in simulating and forecasting severe weather events. All model simulations presented in this article are produced with the resolution of 0.625 degrees in longitude, 0.5 degrees in 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 stage of the fvGCM development. This is a fairly large hurricane with the hurricane eye clearly defined by the vortex structure and by 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, the finite-volume Data Assimilation System (fvDAS), which is based on the fvGCM and is now operational and named as the GEOS-4 DAS.

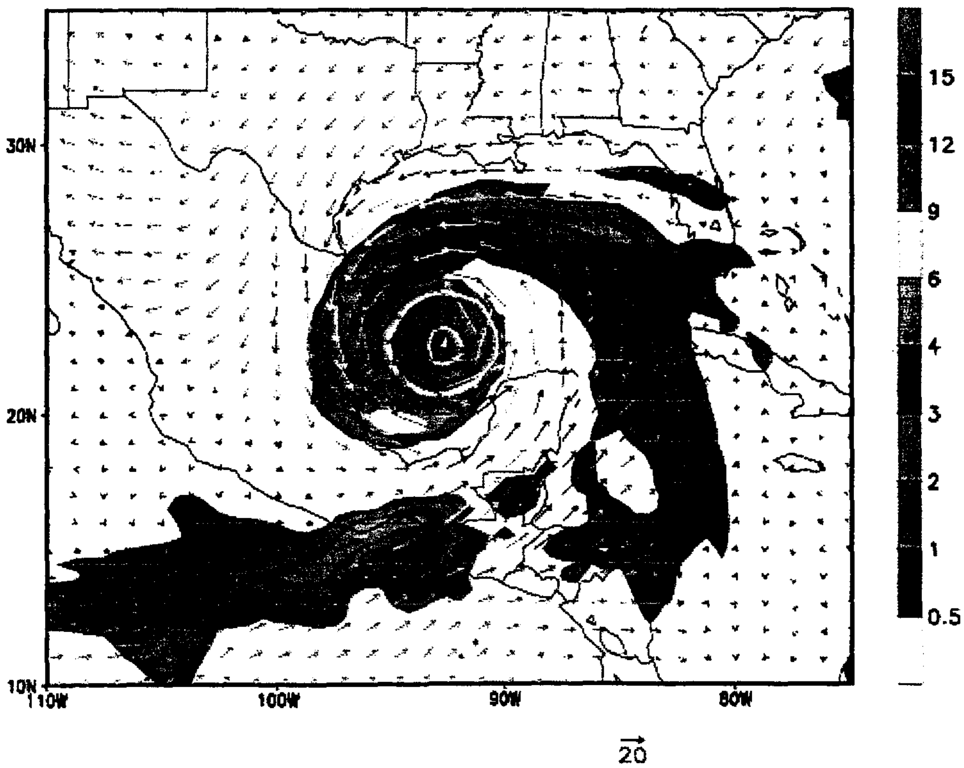

Figure 1: A hurricane simulated by the NASA finite-volume General Circulation Model.
The precipitation rate (mm/hr) is depicted with the color scheme on the right, and the wind (m/s) at the 850mb level is shown with magnitude proportional to the arrow (20 m/s) at the bottom.

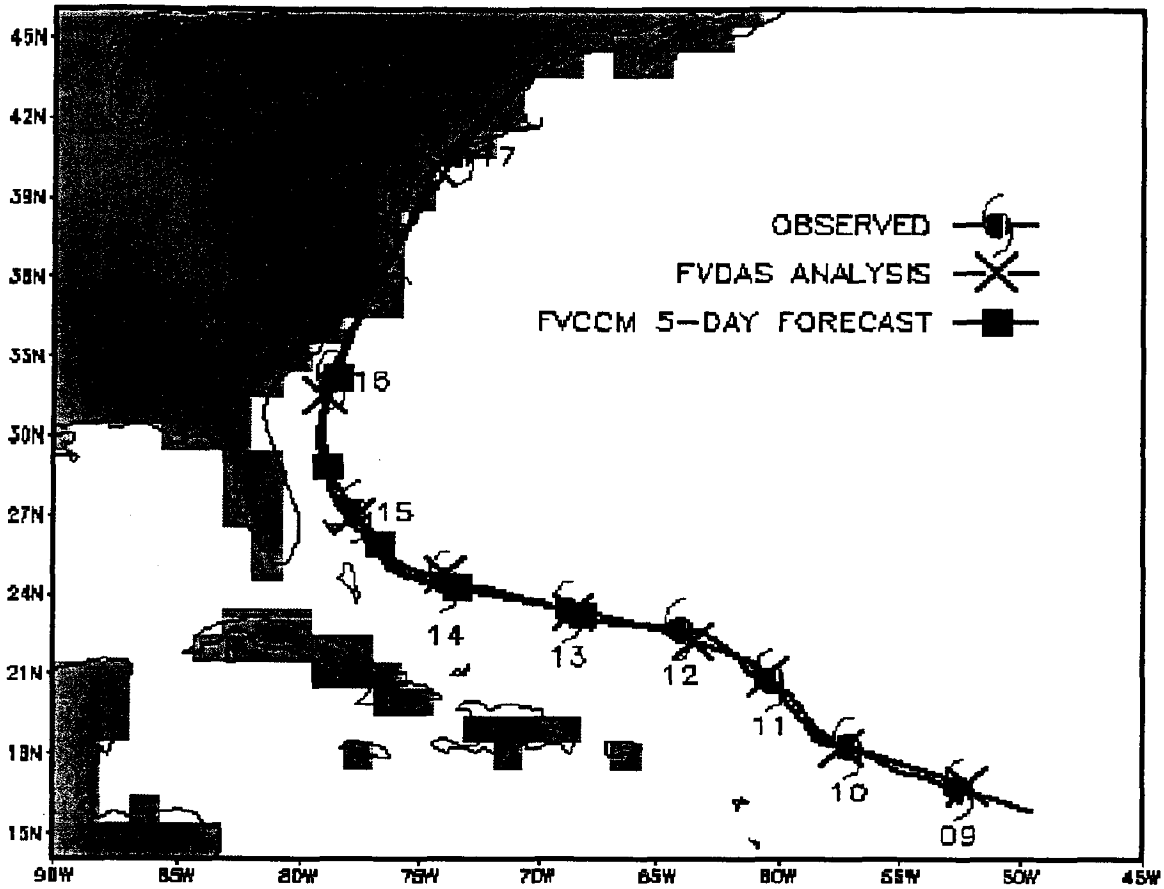

Figure 2: Validation of the NASA finite-volume General Circulation Model (purple squares) and the data assimilation system (blue crosses) with the track of Hurricane Floyd (red spiral spots) observed by the National Hurricane Center in September 1999.

Figure 2 is a validation of the fvGCM and the fvDAS with the track of Hurricane Floyd which occurred in September 1999. The fvGCM was initialized by the fvDAS at 00Z September 12 for a 5-day forecast at the resolution of 0.625x0.5 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 National Hurricane Center.

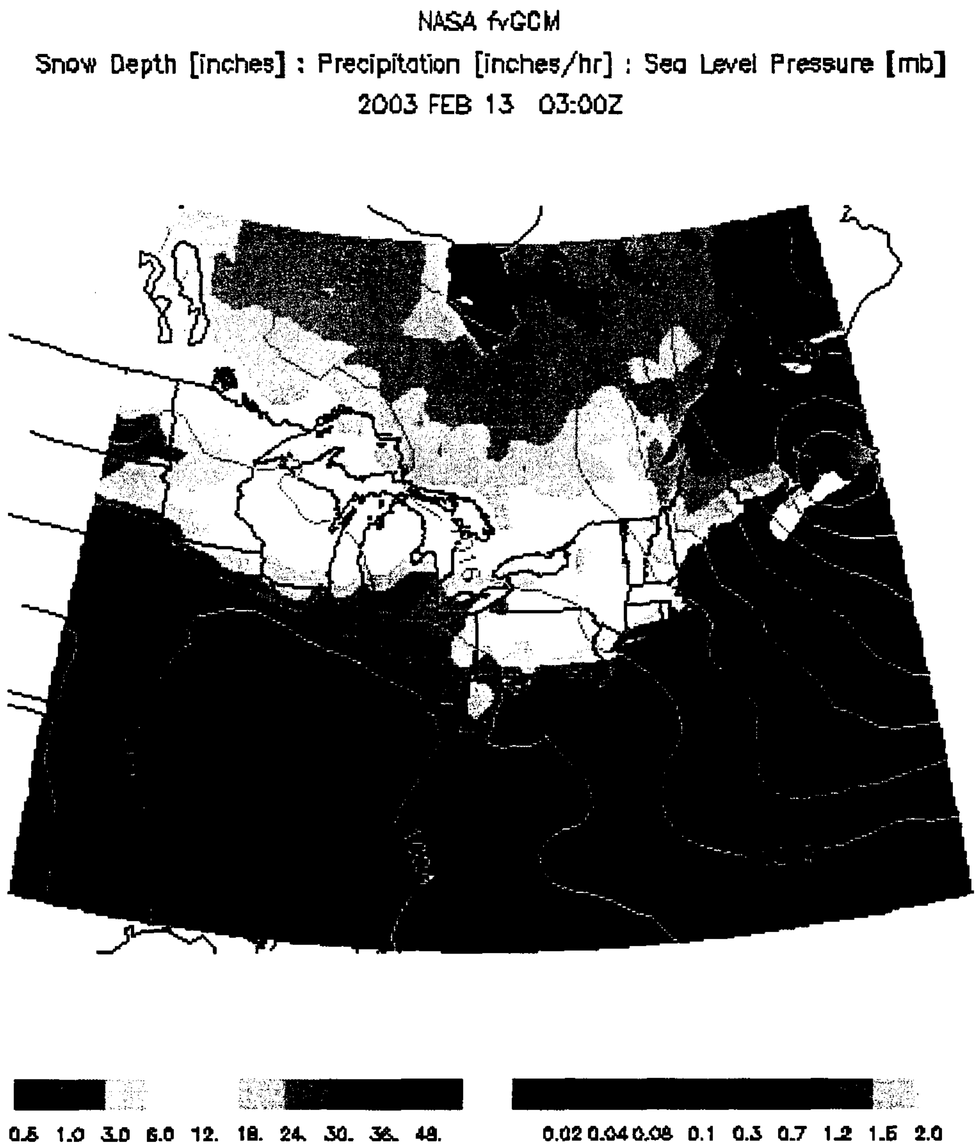

Figure 3: (a) Initial state at 03Z Feb 13, 2003 (sea-level pressure, instant precipitation rate, and the total snow accumulation on the ground) for the snow storm in U.S. east coast in February 2003.

Motivated by the initial success of the fvGCM's forecasts of hurricanes, we started, from the spring of 2002, to evaluate fvGCM's capability in 10-day global weather forecast at real time. The model forecast skill is found to be very competitive with operational centers within the U.S., and we present here a couple of recent severe storms in the following figures. Figure 3 and Figure 4 depicted, at forecast time, the sea-level pressure, instant precipitation rate, and the total snow accumulation on the ground. Figure 4 also includes the winds 10 meters above the ground.

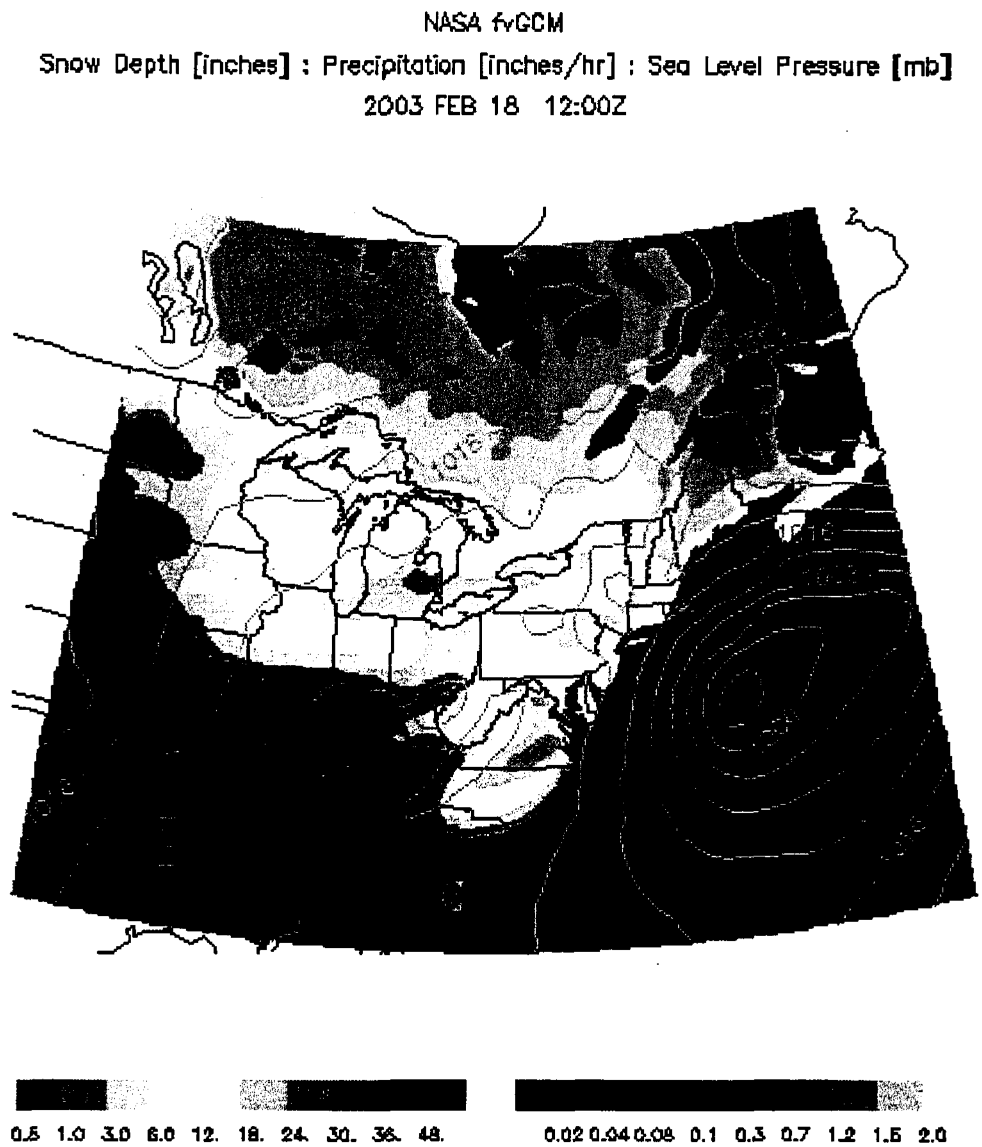

Figure 3: (b) As in (a), but for the 5.5-day forecast at 12Z Feb 18, 2003 .

Figure 3 demonstrates the accuracy of fvGCM in forecasting severe snow storms over the mid-Atlantic states. Part (a) shows the model's initial state at 03Z February 13. Part (b) shows the model prediction after 5.5 days, revealing a 24-inch snow accumulation for the Baltimore-Washington area with amount up to and beyond 30 inches in the mountains of western Maryland and Virginia. The snow accumulation over the entire U.S. east coast is also quite accurate. We, in fact, forecasted this particular "storm of the new century" for the mid-Atlantic states 10 days in advance (in private email exchanges between NASA and NOAA scientists).

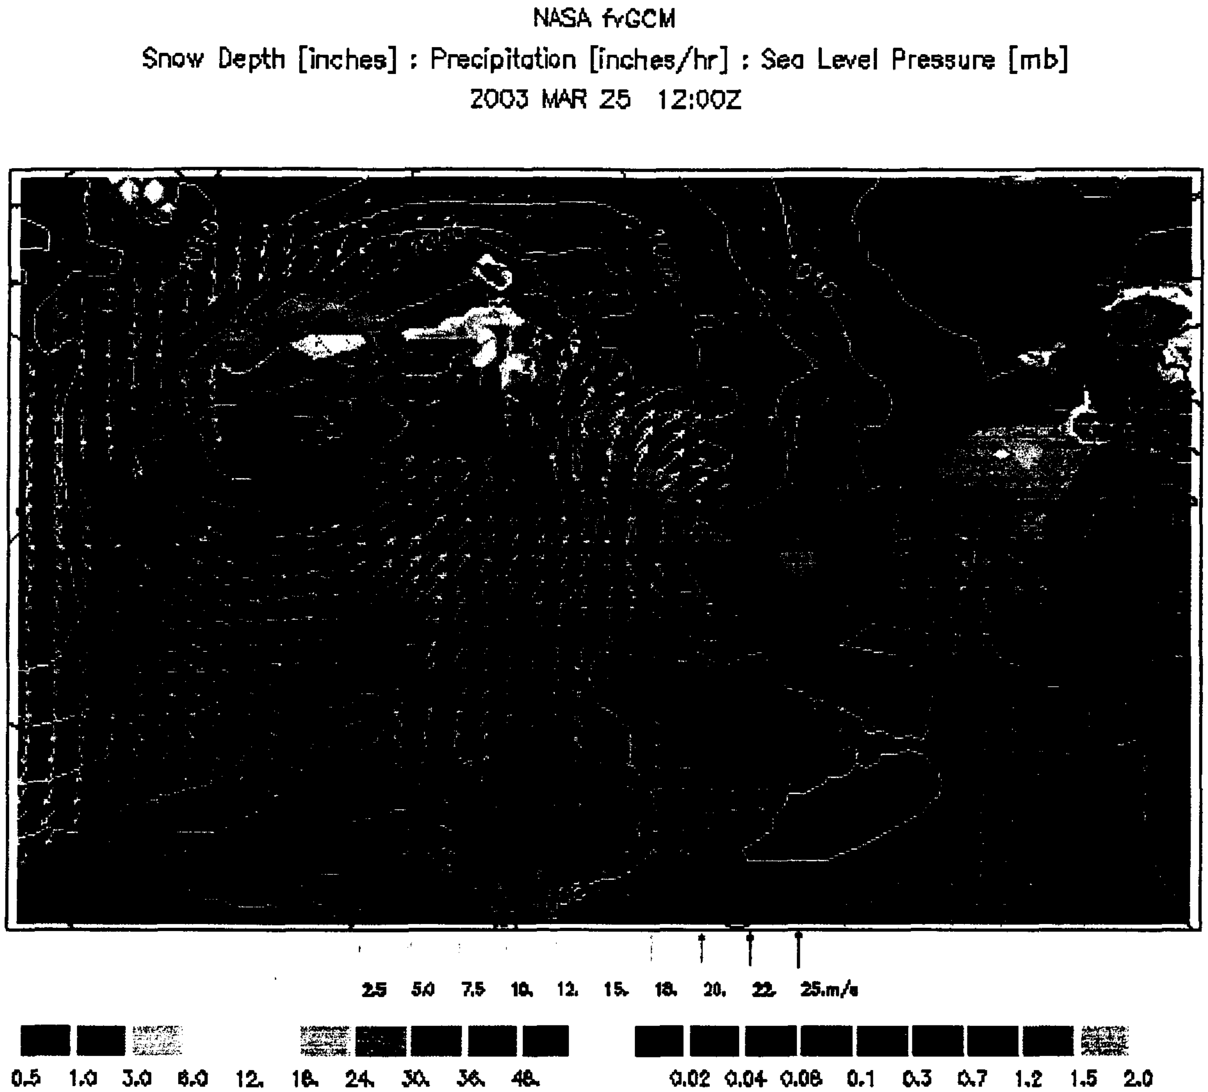

Figure 4: 5.5-day forecast of the severe snow and sand storm in Mideast during the U.S. Operation Iraq Freedom 2003.

Accurate and timely weather forecasts can be useful to military strategies. Immediately after the U.S. Operation Iraqi Freedom in March 2003, we added a window to highlight our real-time global weather forecast in the Mideast. This model forecast was initiated at 00Z March 20, and the severe storm, snow in northern Turkey and sand-storm over the Arabian desert, that happened a week later was successfully predicted by the model. Figure 4 shows a mature severe storm forecasted by the fvGCM on the sixth day of the war in Iraq. The storm approaches from the Mediterranean and Black Seas with strong wind up to 15 m/s (34 miles/hour) in the southwest deserts of Iraq and northern Saudi Arabia. The wind continues to intensify to up to 23 m/s (52 miles/hour) the next day as the storm moves eastward into the land, causing brutal sand-storms in the deserts for 48 hours–exactly as what was described later by the CNN live report. For more details of the NASA fvGCM real-time global weather forecast, the readers are referred to our web site at <*http://polar.gsfc.nasa.gov/sci_research/fvdas/NASCAR_web/nwp*>.

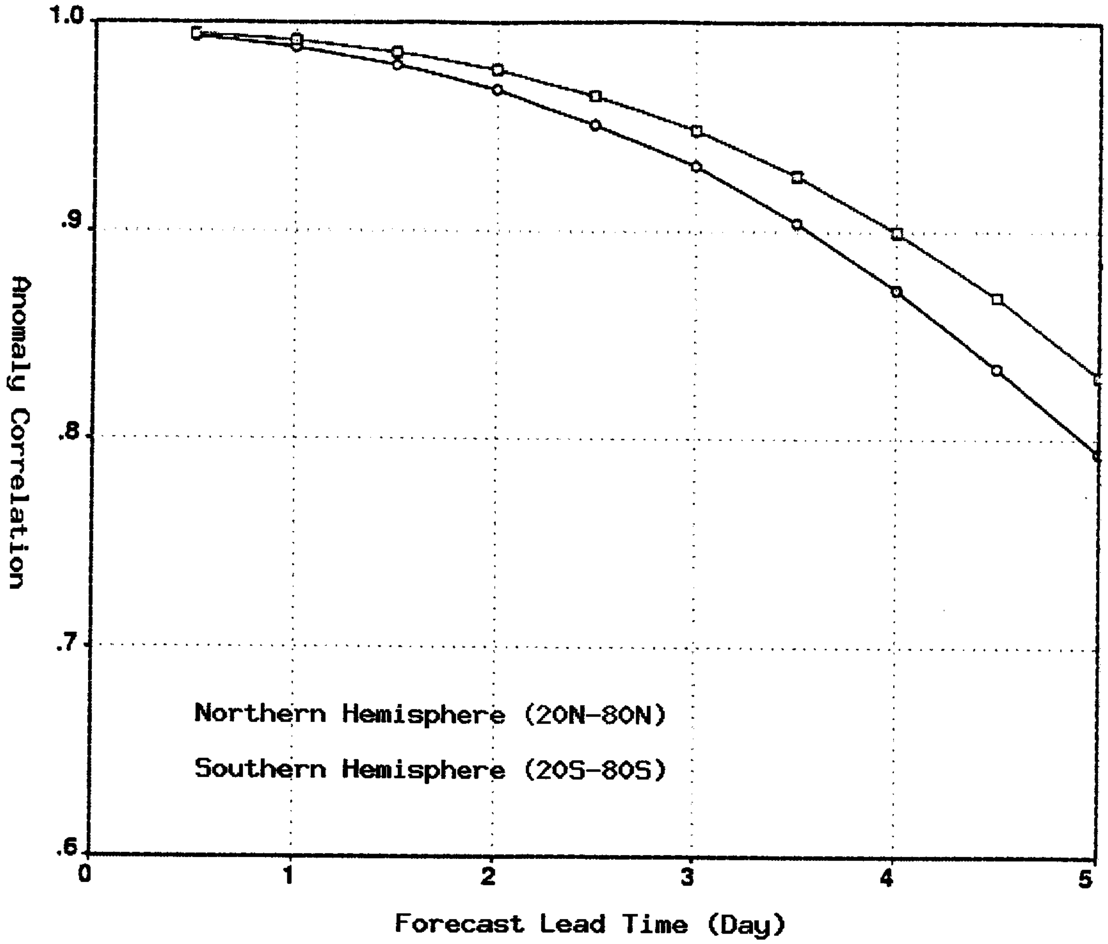

Figure 5: DJF 500mb-height anomaly correlation between the NASA finite-volume General Circulation Model and the analysis by the National Centers for Environmental Prediction during 1 Dec 2002 - 28 Feb 2003. The anomaly correlation in the Northern Hemisphere (winter) is depicted with green rectangles, and in the Southern Hemisphere (summer) with red circles.

To illustrate the fvGCM's general accuracy in global weather forecast, we summarize its forecast skills in Fig. 5 with a 90-day average of *anomaly correlation* – the correlation between model forecast and observation relative to climatology – at the 500mb pressure level which represents the middle troposphere, during the period from December 1, 2002 to February 28, 2003 (DJF). According to the analysis by the National Centers for Environmental Prediction, the fvGCM shows very good forecast skills for up to 5 days in forecast time, with a score of 83% in the Northern Hemisphere (winter) and 79% in the Southern Hemisphere (summer) for the fifth forecast day.

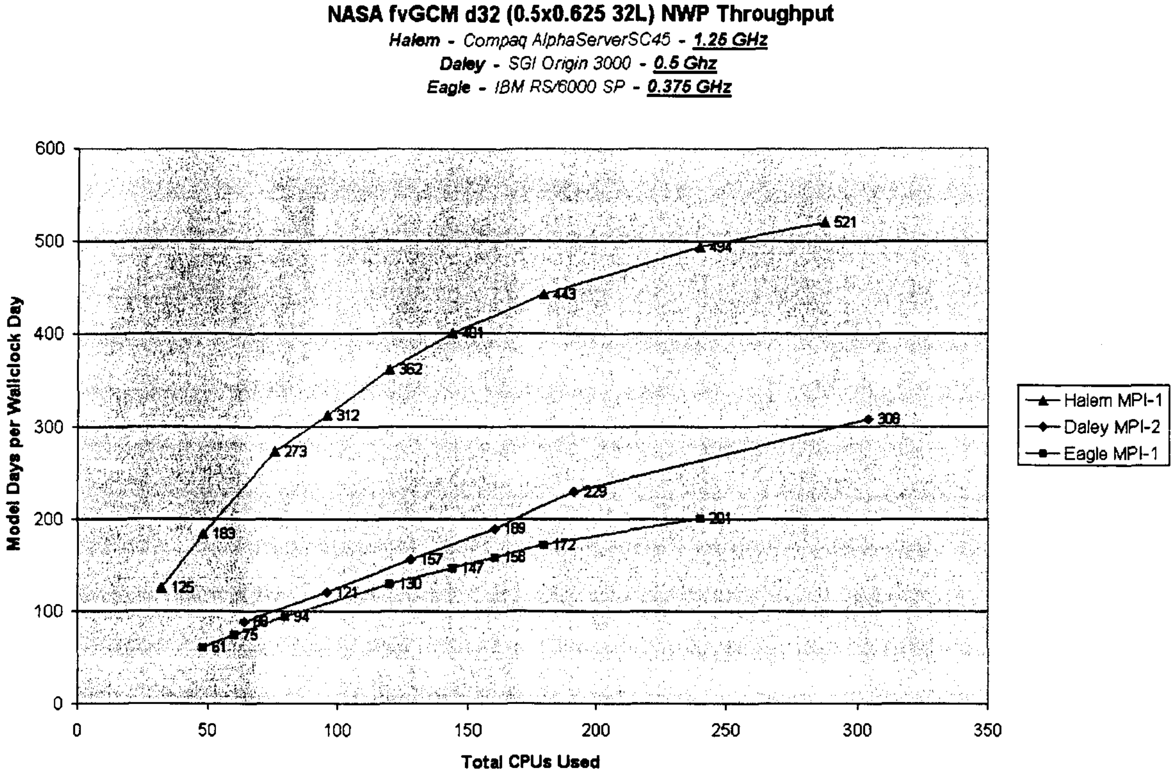

Figure 6: Computational performance of the NASA finite-volume General Circulation Model on various machines, including IBM, SGI and Compaq.

Real-time weather forecasts are time critical, and horizontal resolution 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 additional overhead is required for 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 the optimization of the fvGCM modeling system. Figure 6 shows the fvGCM's computational performance on an IBM RS/6000 SP system named Eagle (magenta 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 increasing number of CPUs is referred to as the ***parallel efficiency*** of the model for the given resolution. Note that the model's throughputs on all three machines increase rapidly with increasing number of CPUs up to about 250, indicating a quite efficient parallel implementation at the operational resolution of 0.625x0.5 degrees. Although the SGI machine Daley appears to have a slightly better scaling in the sense of linear throughput, the processor speed is about 2.5 times faster on Halem than on Daley, hence it is not surprising to see the Compaq machine Halem outperforms the SGI machine Daley in the low- to mid-range CPU counts (32-256).

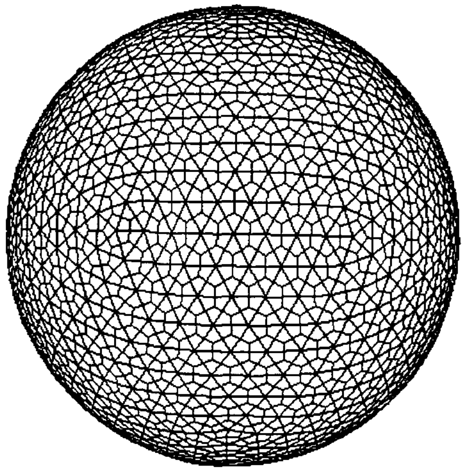

Figure 7: Spring-dynamics icosahedral geodesic grid with gravitational centers for global high-resolution modeling.

3. Future modeling and computing system development

The improvement in parallel efficiency and the advancement in scientific algorithms together can bring as much improvement to the overall computational efficiency as the hardware improvement predicted by the *Moore's law*, which states that the computing power doubles about every 18 months. Although the current fvGCM's parallel efficiency is adequate for today's high-end computers, we would not expect so for future large-scale computing systems that are required to fulfill the nation's goals in weather and climate predictions. As is shown in Fig. 6, the fvGCM's parallel efficiency (slope) on Halem for the current operational resolution (0.625x0.5 degrees) reduces to that of a lower-frequency machine Daley at about 250 CPUs. Today's high-end computers can have as many as 4,000 CPUs in a system, *e.g.*, the Japanese Earth Simulator. To meet, or even to 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 U.S. by 2010. This poses a serious challenge to our ability to improve the model's parallel efficiency.

The current model's limit in parallel efficiency is mainly due to the use of the traditional latitude-longitude grid system where the meridians converge at the computational poles, making a 2D domain decomposition technique less efficient, and even undesirable. To fundamentally resolve the issue of parallel efficiency that is degraded by the presence of computational poles, it seems desirable to abandon the traditional latitude-longitude grid system and seek for a quasi-uniform grid system where there exists no computational poles. For quasi-uniform grid systems, however, it is very difficult to formulate accurate high-order numerical solutions to the physical laws that govern the atmosphere. 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.

Heikes and Randall [7] investigated the convergence problem of numerical solutions on an icosahedral geodesic grid. Tomita et al. [8] 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 (Fig. 7) designed by Tomita et al. for a possible future generation of the fvGCM. We are also looking into the possibility of using a cubed sphere which is obtained by projecting a cube from the Earth center onto the spherical surface. In order to maintain the exceptional quality of the current fvGCM in the future modeling system, we are designing new computational algorithms with the desired properties of local conservation and monotonicity for quasi-uniform grid systems. A prototype of our new finite-volume formulation has been established with the desired numerical properties, and it seems that we will be able to thoroughly resolve the foreseen parallel efficiency issue 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, however, is not sufficient to make realistic weather and climate prediction. A land model, an ocean model and a sea-ice model altogether are needed to provide forcing, in the form of momentum, latent heat, and sensible heat fluxes, to the lower boundary of the atmosphere model. Conversely, 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 the 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 named *Virtual Planet* and to build an exceptionally powerful supercomputer named *Planet Simulator.*

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

- Development of non-hydrostatic 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
- Assimilation of NASA and NOAA high-resolution satellite data

A project this scale would likely require a coordinated national effort involving several agencies, such as DOE, NASA, NOAA and NSF, and research institutions and universities in the United States.

Acknowledgments

We would like to 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, Sept. 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, Oct. 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, July 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., Boston, 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 and Engineering*, vol. 4, no. 5, Sept. 2002, pp. 49-54.

6. Eugenia Kalnay, "Atmospheric modeling, data assimilation and predictability," *Cambridge University Press*, New York, 2003, pp. 341.

7. 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, June 1995, pp. 1862-1880.

8. H. Tomita et al., "Shallow water model on a modified icosahedral geodesic grid by using spring dynamics," *J. Computational Physics*, vol. 174, no. 2, Dec. 2001, pp. 579-613.

BIOGRAPHIES

Dr. 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.

Dr. 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 is a member of the Council of the American Meteorological Society and the Scientific Steering Committee for the U.S. Weather Research Program. He received a BS in aeronautics from St. Louis University, an MS in meteorology and a PhD in meteorology and oceanography from New York University. Contact him at Code 910, NASA Goddard Space Flight Center, Greenbelt, MD 20771; atlas@gmao.gsfc.nasa.gov.

Dr. 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, 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.