

Coupling Advanced Modeling and Visualization to Improve High-Impact Tropical Weather Prediction

To meet the goals of extreme weather event warning, this approach couples a modeling and visualization system that integrates existing NASA technologies and improves the modeling system's parallel scalability to take advantage of petascale supercomputers. It also streamlines the data flow for fast processing and 3D visualizations, and develops visualization modules to fuse NASA satellite data.

The US National Research Council's 2007 Decadal Survey Missions report recommends that "the US government, working in concert with the private sector, academe, the public, and its international partners, should renew its investment in Earth-observing systems and restore its leadership in Earth science and applications." The report includes a top-priority scenario, *Extreme Event Warning*,¹ which focuses on "discovering predictive relationships between meteorological and climatological events and less obvious precursor conditions from massive data sets." To achieve this, we attempted to extend the lead time and reliability of hurricane forecasts (such as track and intensity), which is important for saving lives and mitigating economic damage.

The urgent need for doing this is evidenced by extreme weather events such as Hurricane Katrina in 2005² and Tropical Cyclone Nargis

in 2008,³ which caused tremendous damage and numerous fatalities. Researchers have suggested that large-scale tropical weather systems such as Madden-Julian Oscillations (MJOs),⁴ monsoonal circulations, and tropical easterly waves can regulate tropical cyclone (TC) activity (depending on the location, TCs are referred to by other names, including *hurricanes* in the Atlantic region and *typhoons* in the West Pacific region; they're also variously referred to as tropical storms, cyclonic storms, and tropical depressions). To this end, we sought to improve prediction of these large-scale flows and their impact on TC activities, and thus help extend the lead-time for TC prediction. However, limited computing resources made it challenging to accurately improve these tropical weather systems with traditional global models. Among the major limiting factors in these models are insufficient grid spacing and poor physics parameterizations—such as cumulus parameterizations (CP)—which are designed to "reduce" the model's deficiency with regard to unresolved physical processes because of the coarse-resolution General Circulation Models (GCMs). As the sidebar "Advances in Climate Modeling and Supercomputing" describes, there's great potential now to mitigate these issues. Here, we introduce NASA's supercomputing, concurrent

1521-9615/11/\$26.00 © 2011 IEEE
COPUBLISHED BY THE IEEE CS AND THE AIP

BO-WEN SHEN

University of Maryland, College Park

WEI-KUO TAO AND BRYAN GREEN

NASA

ADVANCES IN CLIMATE MODELING AND SUPERCOMPUTING

In late 2004, NASA's Columbia supercomputer began operation,¹ providing groundbreaking computing power for Earth modeling. Later, the NASA high-end concurrent visualization (CV) system version 1,² in which model outputs are extracted for analysis while the simulation is still running, was developed as a powerful tool for efficiently processing and visualizing massive volumes of high spatial- and temporal-resolution model data.

In late 2008, a new supercomputer, Pleiades, was installed at the NASA Ames Research Center, providing 15 times Columbia's computing power. Enabled by these advanced computational technologies, a high-resolution (approximately 10 km) global model (the finite-volume general circulation model, or fvGCM) was deployed and used to generate remarkable forecasts of intense hurricanes.^{3–5} More importantly, researchers have proposed an innovative approach that uses the fvGCM and a massive number of Goddard Cumulus Ensemble (GCE) models⁶ to overcome the cumulus parameterization (CP) deadlock in GCMs caused by the slow development and inadequate performance of CPs.⁷ This approach is called the *multiscale modeling framework* (MMF) or super-parameterization. In the original implementation, a GCE is run (at a resolution of 4 km, currently) in place of the CP in each of the fvGCM's coarse grids (such as 250 km or 100 km). In the revised implementation with better scalability, these high-resolution GCEs are run collectively as a super-component that spans the same area as the fvGCM. As a result, the MMF has the combined advantages of GCM's global coverage and GCE's sophisticated microphysical processes.

The current MMF consists of a coarse-resolution 100- to 250-km fvGCM and thousands of copies of GCEs; a new version of the MMF will include a high-resolution fvGCM. Both the high-resolution fvGCM and MMF can be run without relying on CPs, and are used to examine the impact of grid-resolved convection and radiation interaction on large-scale tropical systems' simulations.

References

1. R. Biswas et al., "Petascale Computing: Impact on Future NASA Missions," *Petascale Computing: Architectures and Algorithms*, D. Bader, ed., CRC Press, 2007, pp. 29–46.
2. D.B. Ellsworth et al., "Concurrent Visualization in a Production Supercomputing Environment," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, 2006.
3. R. Atlas et al., "Hurricane Forecasting with the High-Resolution NASA Finite Volume General Circulation Model," *Geophysical Research Letters*, vol. 32, no. 3, 2005; doi:10.1029/2004GL021513.
4. B.W. Shen et al., "Hurricane Forecasts with a Global Mesoscale-Resolving Model: Preliminary Results with Hurricane Katrina," *Geophysical Research Letters*, vol. 33, no. 15, 2006; doi:10.1029/2006GL026143.
5. B.W. Shen et al., "Predicting Tropical Cyclogenesis with a Global Mesoscale Model: Hierarchical Multiscale Interactions During the Formation of Tropical Cyclone Nargis (2008)," *J. Geophysical Research*, vol. 115, 2010; doi:10.1029/2009JD013140.
6. W.K. Tao and J. Simpson, "The Goddard Cumulus Ensemble Model. Part I: Model Description," *Terrestrial, Atmospheric and Oceanic Sciences*, vol. 4, 1993, pp. 35–71.
7. W.K. Tao et al., "A Goddard Multi-Scale Modeling System with Unified Physics," *World Climate Research Programme/Global Energy and Water Cycle Experiment (WCRP/GEWEX) Newsletter*, vol. 18, no. 1, 2008, pp. 6–8.

visualization (CV), and global modeling technologies and discuss how we built the Coupled Advanced Multiscale Modeling and CV Systems (CAMVis) by

- integrating existing NASA technologies, such as the NASA multiscale modeling system, the Goddard Cumulus Ensemble (GCE) model, the finite-volume GCMs, and CV systems;
- improving the parallel scalability of the coupled multiscale modeling system to take full advantage of petascale supercomputers;
- significantly streamlining data flow for fast processing and 3D visualizations; and
- developing visualization modules for the fusion of NASA satellite data—including precipitation from the Tropical Rainfall Measuring Mission (TRMM) and surface winds from the Quick Scatterometer (QuikSCAT).

We also describe our project's progress and how CAMVis can help provide insightful understanding of multiple physical processes and their multiscale interactions with improved short-term (approximately five- to 10-day) forecasts of TCs and extended-range (30-day) simulations of large-scale MJOs.

Supercomputing and Modeling Technology at NASA

In late 2004, the Columbia supercomputer⁵ came into operation at NASA's Ames Research Center. It consists of 20 512-CPU nodes, giving it 10,240 CPUs and 20 terabytes of memory (see Figure 1). Columbia achieved a performance of 51.9 trillion floating-point operations per second (Tflops/s) with the Linpack benchmark. These large-scale computing capabilities let researchers solve complex problems using large-scale modeling systems.^{2,6}



(a)



(b)

Figure 1. NASA supercomputers. (a) The Columbia supercomputer with (in late 2004) 20 SGI Altix superclusters, 10,240 Intel Itanium II CPUs, and 20 Tbytes total memory. (b) The Pleiades supercomputer with 111,104 cores (with Xeon, Nehalem, and Westmere processors), more than 120 Tbytes of memory, and more than 6.9 petabytes of disk space.

In late 2008, the Pleiades supercomputer, an SGI Altix ICE system with a peak performance of 609 Tflops/s, was built as one of the most powerful general-purpose supercomputers. Recently, Pleiades has been upgraded to have 111,104 cores with a Linpack performance of 1,088 Tflops, 185 Tbytes memory, and 6.9-petabyte disk space. This new system, which provides more than 15 times the computing power of Columbia, is expected to speed up scientific discovery at an unprecedented pace.

Concurrent Visualization System


It's well known that the substantial increase in data volume produced by high-resolution Earth modeling systems poses a great challenge to stage, handle, and manage these model outputs and compare them with satellite data. We believe that efficiently handling these massive data sets, from terabytes for short-term runs to petabytes for long-term runs, requires an innovative thought process and approach. CV is a technique that could achieve this goal and has met with great success in visualizing high-volume data.^{7,8} In CV, a simulation code is instrumented such that its data can be extracted for analysis while the simulation is running without having to write the data to disk. By avoiding filesystem I/O and storage costs, CV provides much higher temporal resolution than is possible with traditional post-processing, enabling every time step of a very high-resolution simulation to be captured for analysis. The other main benefit of CV is that it provides a view of a simulation in progress, which can be useful for application monitoring or steering. This can help detect serious job failure and avoid wasting system resources.

In 2005, CV technology was first developed and integrated into the high-resolution finite-volume


general circulation model (fvGCM) on the original hyperwall system (49 screens). A 3×3 screen "mini-hyperwall" was used for looping the resulting movies. Recently, we deployed an improved CV system (version 2; see Figure 2a) that consists of a front-end system for data extraction ("coalescer"), a middle-layer system for data handling and data rendering, and a back-end system for data display. "Extractions" include domain slices or subvolumes, cutting planes, isosurfaces, streamlines, and other feature-extraction products. An extract's size can vary widely, depending on what features are captured, while a movie's size is determined only by the image resolution and the compression level achieved during encoding. The great advantage of extracts is that they represent an intermediate data product, which can be loaded into a viewer at a later time for interactive analysis.

As of June 2008, NASA's 128-screen hyperwall-2, capable of rendering one-quarter-billion pixel graphics, was built at NASA's Ames Research Center as one of the world's highest-resolution scientific visualization and data exploration systems. Compared to the original 49 screens and 100 Base-T interconnect, the hyperwall-2 has 128 screens with modern graphics cards, an InfiniBand interconnect, and is fully integrated into the NASA supercomputing environment. The hyperwall-2's 1,024-CPU cores and 475 Tbytes of fast disk provide an excellent environment for parallel feature extraction and extract storage. In addition, the hyperwall-2's high-speed interconnect makes fully 3D CV possible.

To efficiently exchange data between the computing and visualization nodes, we've implemented the *M-on-N* configuration for the CV pipeline, as Figure 2b shows. The *M-on-N* configuration



(a)



(b)

Figure 2. Architecture of the new concurrent visualization (CV) system (version 2) with parallel data transfer configuration. (a) In addition to computing nodes (“Pleiades”), the system consists of a front-end system (“coalescer”) and the hyperwall-2 128-node, 8-core/node rendering cluster. Rounded rectangles indicate systems and regular rectangles indicate processes. These systems are used for data extraction, handling, and visualization and for MPEG image production and visualization display. (b) The M -on- N configuration for parallel data communications between the M computing nodes and N visualization nodes in the CV.⁸

allows different domain decomposition within computing and visualization nodes. After decomposition, each portion of the entire domain within a computing (visualization) node is referred to as a *subdomain* (subregion). The boxes on the left represent the multiple (M) MPI processes of a fvGCM job, with each one responsible for simulation within a subdomain. At startup, each MPI process in the fvGCM job creates a connection with the Infiniband remote data memory access (RDMA) protocol to one of the N MPI processes spawned from the hyperwall job, using an M -on- N

mapping where $M \geq N$. At the end of each time step, each computing node’s raw output is transferred directly via its RDMA connection to its corresponding peer process on the specific hyperwall node. The hyperwall job then performs feature extraction and “sort-last” rendering in parallel, wherein each child MPI renders an image from its own data. These “individual” images from all of the child MPI processes are then sorted and composited into a complete image in Portable Pixel Map (PPM) format, which can be passed to an encoder for movie generation. Finally, a series




Figure 3. The NASA global multiscale modeling framework (MMF). The system consists of the global model (the finite-volume general circulation model, or fvGCM) and thousands of copies of cloud models (such as the Goddard Cumulus Ensemble, or GCE). (a) Visualization of a five-day, low-level wind simulation with the global model during the Nargis period (Figure 5 shows details on this). (b) 3D visualization of cloud water and ice are depicted in white, with nine copies of a cloud model in a 6-degree by 7.5-degree domain, thus demanding 13,104 copies of the cloud model in the global environment.

of these complete images are converted to JPEG, which can be easily delivered to a Web server for display.

To maximize a single simulation run's results, multiple products are usually generated, representing various fields and regions of interest, and numerous feature-extraction and visualization techniques. As part of the CV pipeline, the resulting animations are streamed as they're generated to remote displays at the principal scientists' facilities. When time-stamped outputs arrive from the computing nodes, each visualization node sequentially computes all requested visualizations, producing one image per visualization request. The node that assembles the final composite image is assigned in a round-robin fashion, so that the encoders are spread out across the cluster.

Depending on the visualization produced, additional data exchange might occur within the hyperwall (visualization) nodes. Visualizations such as scalar volume rendering, cutting planes, and iso-surfaces are easily implemented within the "sort-last" renderer—where each processor has a subset of the overall scene geometry that it uses to produce an image with both color and depth information—with only ghost-cell exchange needed for the subdomain boundaries. (In parallel computing, ghost cells are points outside the target domain's boundary that are needed for computing, including rendering). However, for performing vector visualization techniques within the subregions—which might cover several subdomains on different nodes—it might be convenient to fully reconstruct a target subregion on each of the visualization nodes.

Multiscale Modeling System

Researchers at NASA's Goddard Space Flight Center (GSFC) successfully developed the first version of the multiscale modeling system with unified physics⁶ and deployed it on the Columbia supercomputer. It has since been ported and tested on the Pleiades supercomputer.

As Figure 3 shows, the system consists of the fvGCM^{2,9,10} at a coarser (100- to 250-km) resolution and thousands of copies of a cloud model (such as GCE¹¹) at a 4-km or finer resolution. With the current model configurations, 13,104 GCEs are run collectively to explicitly simulate cloud processes in the global environment, providing cloud feedbacks to the atmospheric state in the fvGCM. The high-resolution fvGCM was first deployed on the Columbia supercomputer (and later on Pleiades), producing remarkable forecasts of intense hurricanes in 2004 and 2005.^{2,3} Both the multiscale modeling framework (MMF) and high-resolution fvGCM, which can be run with no dependence on CPs, are powerful tools for examining and understanding the impact of grid-resolved convections.

Computational Enhancement

Over the past few years, a single-program multiple-data (SPMD) parallelism has been separately implemented in both the fvGCM and GCE with good parallel efficiency.^{12,13} However, the key for improving the overall performance is to increase the copies of the GCEs to be run in parallel, because at runtime, 95 percent or more of the total wall-time for running the MMF is spent on the GCEs' multiple copies. Thus, wall time could be significantly reduced by efficiently distributing

the numerous GCEs over a massive number of processors on a supercomputer. However, the original implementation—in which each of the 13,104 GCEs is embedded on a grid point in the fvGCM—has very limited parallel scalability, with the total number of CPUs limited to 30.

As we describe later, we propose a different strategic approach to overcome this difficulty that couples the fvGCM and GCEs. From a computational perspective, the concept of “embedded GCEs” should be completely forgotten, as it restricts the view on the fvGCM’s data parallelism. Instead, the 13,104 GCEs should be viewed as a meta-global GCE (mgGCE) in a meta-gridpoint system. With this concept in mind, each of the MMF’s two distinct “components”—the fvGCM and mgGCE—could have its own scaling properties. Because most of the wall time in MMF runs is spent on the GCEs, a scalable mgGCE could substantially reduce wall time. In addition, it becomes feasible to implement a multiple-programs, multiple-data (MPMD) parallelism for the MMF with the mgGCE and the fvGCM.

Currently, the coarse-resolution fvGCM is running with 1D domain decomposition because it costs a small percentage of the wall time in MMF runs. Briefly, the technical approaches for this implementation with 2D domain decomposition in the mgGCE are as follows:

1. A master process allocates a shared-memory arena for data redistribution between the fvGCM and mgGCE by calling the Unix `mmap` function.
2. The master process spawns multiple (parent) processes with a 1D domain decomposition in the y direction through a series of Unix `fork` system calls.
3. Each of these parent processes then forks several child processes with another 1D domain decomposition along the x direction in the mgGCE.
4. Data gathering in the fvGCM (mgGCE) is done along the y direction (along the x direction and then the y direction).
5. Synchronization is implemented with the atomic `__sync_add_and_fetch` function call on the Columbia supercomputer.

Although steps 1, 2, and 5 were previously used in multiple-level parallelism, this methodology has been extended to the multicomponent system (namely, the fvGCM and mgGCE¹⁴). Highly promising scalability—up to 364 CPUs—has been achieved with preliminary benchmarks, which

show a speedup of 3.93, 7.28, and 12.43 by increasing the number of CPUs from 30 to 91, 182, and 364, respectively.¹⁴ This encouraging speedup is achieved largely because the current mgGCE, in which each GCE running with periodic lateral boundary conditions, has no ghost cells among different GCEs. We’re integrating this new MMF (with a scalable mgGCE) and CV—that is, the CAMVis—on the Pleiades and hyperwall-2. We’re also working on further enhancements, including scalability and functionality.

Scientific Applications

We now discuss CAMVis system visualizations with improved convective/cloud processes to illustrate

- improved short-term (approximately five- to seven-day) predictions for the formation of twin TCs associated with a large-scale MJO;

Researchers have documented that the nearly simultaneous formation of two tropical cyclones straddling the equator at low latitudes occasionally can occur in the Indian Ocean and West Pacific Ocean.

- insightful visualizations of multiple processes and their scale interactions, which led to the formation of TC Nargis (2008);
- improved extended-range (approximately 30-day) simulations of a large-scale MJO; and
- comparisons of model simulations and satellite measurements at comparable resolutions.

We selected these cases to provide a detailed (zoomed-in) view on hurricane physical processes and an integrative (zoomed-out) view on its interactions with environmental conditions.

Simulating Twin Cyclone Formation Associated with an MJO

Researchers have documented that the nearly simultaneous formation of two TCs straddling the equator at low latitudes occasionally can occur in the Indian Ocean and West Pacific Ocean. These TCs are called “twins” as they’re nearly symmetric with respect to the equator. Previous studies showed that this twin TC activity can be modulated by the large-scale MJO.




Figure 4. Predictions for twin tropical cyclone formations in the Indian Ocean. (a) Madden-Julian Oscillation (MJO)-organized convection over the Indian Ocean at 0630 UTC 1 May 2002. When the MJO moved eastward, two pairs of twin tropical cyclones—TC 01A, Kesiny, TC 02B, and Errol—appeared sequentially on (b) 6 May and (c) 9 May. Two of the TCs (01A and 02B) had counterclockwise circulations and appeared in the Northern Hemisphere, while the other two (Kesiny and Errol) had clockwise circulations and appeared in the Southern Hemisphere. (d) A four-day forecast of total precipitable water, showing realistic simulations of TC formation and movement (as we'll detail in a forthcoming article, we verified the simulated tracks and precipitations in 10-day runs against observations).

In May 2002, for example, large-scale organized convection associated with an MJO event was observed in the Indian Ocean (see Figure 4a). While the MJO was continuously progressing eastward, two pairs of twin TCs appeared (Figures 4b and 4c). To capture the genesis, a 10-day forecast was initialized at 0000 UTC 6 May (Figure 4d). The genesis and movement of three of these TCs (02B, 01A, and Kesiny) were simulated realistically; however, for the southern entity of the second pair of twin TCs (Errol), only less-organized convection was simulated.

3D Visualization of Nargis

In 2008, Nargis—a severe cyclonic storm that's the deadliest-named TC in the North Indian Ocean Basin—caused more than 133,000 fatalities and \$10 billion in damage. An increased lead time in the prediction of Nargis would have increased

the warning time and might therefore have saved lives and reduced economic damage. Global high-resolution simulations using real data³ showed that the initial formation and intensity variations of Nargis could be realistically predicted with position errors of 200 km at a lead time of up to five days. Experiments also suggested that the accurate representation of environmental flows such as a westerly wind burst associated with an MJO is important for predicting the formation of this kind of TC. As we discuss later, providing a simplified view of these multiple processes using 3D visualizations could prove quite powerful.

In contrast to the twin TC case, simulating and understanding processes for the developing TC Nargis and the nondeveloping vortex (the counterpart to Nargis in the Southern Hemisphere) are equally important. To illustrate this, we produced




Figure 5. Realistic seven-day simulations of the formation and initial intensification of TC Nargis (2008) initialized at 0000 UTC 22 April 2008, showing streamlines at different levels. Low-level winds are blue and upper-level winds are red. (a) Formation of a pair of low-level vortices (V) at the 84th hour of simulation. (b) Intensification of the northern vortex (left). (c) Formation of Nargis associated with the enhancement of the northern vortex. (d) Intensification of Nargis associated with upper-level outflow and moist processes, indicated by the enhanced upper-level outflow circulation. Approaching upper-level easterly winds (E) increase the vertical wind shear, suppressing the enhancement of the southern vortex (on the right).

a set of 3D, high-temporal-resolution animations with CAMvis. Figure 5 shows snapshots of streamline visualization at different vertical levels. Low-level winds are shown in blue, and upper-level winds in red. In Figure 5a, ending at 1200 UTC 25 April 2008, a pair of low-level vortices (V) appeared in the Northern and Southern Hemispheres, showing the potential for the formation of a pair of twin TCs. As time progressed, the (low-level) westerly wind belt/burst (W) moved northward, enhancing the horizontal wind shear and therefore intensifying the northern vortex into Nargis (Figures 5b and 5c). With other favorable conditions, including good upper-level outflow, Nargis continued to intensify (Figure 5d). In contrast, at 0000 UTC April 26 (Figure 5b), upper-level easterly winds (labeled “E”), which moved over the top of the southern vortex, increased the vertical wind shear and therefore suppressed the enhancement of

the southern vortex (Figure 5b). Other unfavorable factors (such as the proximity to the equator) also contributed to the lack of TC formation in the Southern Hemisphere during this period.

MJO Simulations with the MMF

In the previous section, we showed the TC formation associated with MJOs; we now discuss the model’s performance for simulating an MJO. Accurately predicting tropical activity at sub-seasonal scales (approximately 30 days) is crucial for extending numerical weather prediction beyond two weeks, and accurate forecasting of an MJO is among the challenges. With a 45- to 60-day time scale, eastward-propagating MJOs, which are typically characterized by deep convection originating over the Indian Ocean, have one of the most prominent large-scale features of the tropical general circulation. The MMF




Figure 6. A 30-day simulation of an MJO initialized at 0000 UTC 13 December 2006, as shown by the 200-hecto Pascals (hPa) velocity potential. This simulation captures several major features usually associated with an MJO: (a) Day 0 at the initial time. (b) Day 3: initiation of large-scale organized convection in the Indian Ocean. (c) Day 8: intensification. (d) Day 13: slow propagation (prior to reaching the Maritime continent). (e) Day 18: fast propagation. (f) Day 23: weakening. However, the MJO also produces stronger vertical motion than observations.

provides an innovative approach to investigating the multiple processes and multiscale interactions that are important for improving MJO simulations.

Figure 6 shows the 30-day simulation of an MJO event initialized at 0000 UTC 13 December 2006, illustrating that the MJO's life cycle is successfully captured and therefore that the model has the potential to help us examine an MJO's impact on climate simulations.

Comparing Model Simulations and Satellite Measurements

Precipitation is a good indicator of the energy source of an intensifying TC and of low-level wind speeds for measuring TC intensity. Data fusion of NASA's Tropical Rainfall Measurement Mission (TRMM) precipitation and QuikSCAT winds into the CAMVis system is therefore valuable for comparing various high-resolution model simulations with various satellite measurements.

We recently developed data conversion and visualization modules for this purpose. Figures 7a and 7b show TRMM precipitation and QuikSCAT winds during the lifetime of Nargis.

Given the new data fusion capability, we compared the QuikSCAT winds for Nargis with high-resolution model simulations. Our goal was to assess the data-consistent accuracy in the representation of mesoscale vortex circulation and thus improve formation prediction. The assurance of data continuity (or consistency) is important for accurately tracing a TC's movement or identifying its formation. As the zoomed-in panels of Figure 8 show, the changes of vortex structure aren't smooth (see, for example, the less realistic vortex in the lower left visualization in Figure 8). This suggests the potential for rainfall contamination in the derived wind distributions, which might impact the detection of a TC's formation.




Figure 7. Initial implementation of a visualization module into the Coupled Advanced Multiscale Modeling and CV Systems (CAMVis) information system. The module includes (a) a data converter for Tropical Rainfall Measurement Mission (TRMM) satellite-derived precipitation and (b) a vector plotter of Quick Scatterometer (QuikSCAT) winds.




Figure 8. Vector visualizations of NASA QuikSCAT winds during the initial formation and intensification of Nargis from 1200 UTC 26 April 2009 at a time interval of 12 hours. Zoomed-in windows are used to track the evolution of the mesoscale vortex with a closed circulation at 04/26/12z, 04/28/00z, and 04/28/12z, respectively.

To support NASA missions and reduce the time to scientific discovery, we propose to seamlessly integrate the NASA advanced modeling and supercomputing technologies. Our plan is to improve the CAMVis system's performance by taking full advantage of Pleiades' computing power; to

improve the simulations of cloud processes with 3D GCEs; and to implement and test more sophisticated cloud schemes. We'll improve the coupled system to address the interactions of clouds, radiation, and aerosols, to advance our understanding of the detailed 3D structure of these fields and investigate their impact on tropical weather

prediction by comparing these high-resolution simulations with NASA high-resolution satellite observations. These satellites include current missions, such as TRMM and QuikSCAT, as well as future missions such as Global Precipitation Measurement, Aerosol-Cloud Ecosystems, and the 3D-Winds missions described in the 2007 Decadal Survey Report.

We successfully developed and tested each of the CAMVis system's individual components on the Columbia supercomputer (including the high-resolution fvGCM,² GCE and MMF version 1,⁶ and CV version 1⁷). We recently finished deploying new versions of individual components—including a performance-enhanced MMF¹⁴ and CV version 2⁸—and the initial CAMVis system on NASA's Pleiades supercomputer. Because the multiple-scale modeling system can simulate weather and climate at high spatial and temporal resolution, coupling these modeling and CV systems can help process massive volumes of output efficiently and provide insightful understanding of the complicated physical processes. The CV system is equipped with the 128-screen hyperwall-2 and is connected via high-speed InfiniBand to the Pleiades supercomputer. The CV system has several key benefits, including that it lets us

- monitor system runtime status and thereby detect serious failures that could waste system resources;
- use much higher temporal resolution, as it largely obviates I/O and storage space requirements; and
- concurrently visualize complicated physical processes with 3D visualizations.

We'll continue to improve the CAMVis system's accuracy and computational performance. Although the high-resolution fvGCM can generate five-day forecasts in real time, our short-term goal is to scale the CAMVis—and the MMF and mgGCE—up to perhaps 3,000 CPUs to finish five-day real-time forecasts. A long-term goal is to take full advantage of Pleiades to scale the model (up to 13,104 or higher cores, for example) and thus improve long-term climate simulations. Our vision is that the ultimate CAMVis system will enable researchers, policy and decision makers, and educators to monitor global model simulations at a wide range of spatial and temporal resolutions in real time.

Acknowledgments

We thank the reviewers for their valuable suggestions, which have substantially improved this article. We're

also grateful to the following organizations for their support: the NASA Earth Science Technology Office; the Advanced Information Systems Technology Program; the US National Science Foundation Science and Technology Center; the NASA Modeling, Analysis Prediction Program; the Energy and Water Cycle Study; the NASA High-End Computing Program; the NASA Advanced Supercomputing facility at Ames Research Center; and the NASA Center for Computational Science at Goddard Space Flight Center. Finally, we thank Steve Lang for proofreading this manuscript.

References

1. US National Research Council Decadal Survey, *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond*, The Nat'l Academies Press, 2007.
2. B.W. Shen et al., "Hurricane Forecasts with a Global Mesoscale-Resolving Model: Preliminary Results with Hurricane Katrina," *Geophysical Research Letters*, vol. 33, no. 15, 2006; doi:10.1029/2006GL026143.
3. B.W. Shen et al., "Predicting Tropical Cyclogenesis with a Global Mesoscale Model: Hierarchical Multi-scale Interactions During the Formation of Tropical Cyclone Nargis (2008)," *J. Geophysical Research*, vol. 115, 2010; doi:10.1029/2009JD013140.
4. R.A. Madden and P.R. Julian, "Detection of a 40–50 Day Oscillation in the Zonal Wind in the Tropical Pacific," *J. Atmospheric Science*, vol. 28, no. 5, 1971, pp. 702–708.
5. R. Biswas et al., "Petascale Computing: Impact on Future NASA Missions," *Petascale Computing: Architectures and Algorithms*, D. Bader, ed., CRC Press, 2007, pp. 29–46.
6. W.K. Tao et al., "A Goddard Multi-Scale Modeling System with Unified Physics," *World Climate Research Programme/Global Energy and Water Cycle Experiment (WCRP/GEWEX) Newsletter*, vol. 18, no. 1, 2008, pp. 6–8.
7. D.B. Ellsworth et al., "Concurrent Visualization in a Production Supercomputing Environment," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, 2006, pp. 997–1004.
8. B. Green, C. Henze, and B.-W. Shen, "Development of a Scalable Concurrent Visualization Approach for High Temporal- and Spatial-Resolution Models," *Eos Trans. Am. Geophysical Union*, vol. 91, no. 26, 2010; http://atmospheres.gsfc.nasa.gov/cloud_modeling/docs/Green_Henze_Shen_2010.pdf.
9. S.J. Lin, "A 'Vertically Lagrangian' Finite-Volume Dynamical Core for Global Models," *Monthly Weather Review*, vol. 132, no. 10, 2004, pp. 2293–2307.
10. R. Atlas et al., "Hurricane Forecasting with the High-Resolution NASA Finite Volume General Circulation Model," *Geophysical Research Letters*, vol. 32, no. 3, 2005; doi:10.1029/2004GL021513.


11. W.K. Tao and J. Simpson, "The Goddard Cumulus Ensemble Model. Part I: Model Description," *Terrestrial, Atmospheric and Oceanic Sciences*, vol. 4, 1993, pp. 35–71.
12. J.-M. Juang et al., "Parallelization of NASA Goddard Cloud Ensemble Model for Massively Parallel Computing," *Terrestrial, Atmospheric and Oceanic Sciences*, vol. 18, no. 3, 2007, pp. 593–622.
13. W. Putman, S.J. Lin, and B.W. Shen, "Cross-Platform Performance of a Portable Communication Module and the NASA Finite Volume General Circulation Model," *Int'l J. High Performance Computing Applications*, vol. 19, no. 3, 2005, pp. 213–223.
14. B.W. Shen et al., "Scalability Improvements in the NASA Goddard Multiscale Multicomponent Modeling Framework for Tropical Cyclone Climate Studies," *Proc. Int'l Conf. High-Performance Computing in Asia-Pacific Region*, 2009, pp. 249–256; http://atmospheres.gsfc.nasa.gov/cloud_modeling/docs/Shen_HPC_ASIA_2009_Proceedings.pdf.

Bo-Wen Shen is a research scientist at the University of Maryland, College Park, and the NASA Goddard Space Flight Center, and principal investigator of the CAMVis project, where he's currently working to improve the multiscale interactions of large-scale tropical waves and their modulation on the tropical cyclone activities. His research interests include high-resolution global and regional modeling, high-end

computing, and numerical hurricane and weather prediction. Shen has a PhD in atmospheric sciences from North Carolina State University. He is a member of American Geophysical Union. Contact him at bowen.shen@gmail.com.

Wei-Kuo Tao is a senior research meteorologist at the NASA Goddard Space Flight Center and is the primary developer of the GCE model, the principal investigator on the development of the fvMMF, and the leader of Goddard Mesoscale Modeling and Dynamic Group. His research interests include cloud physics and modeling mesoscale convective systems. Tao has a PhD in atmospheric sciences from the University of Illinois. He is a fellow of the American Meteorological Society and Royal Meteorological Society. Contact him at wei-kuo.tao-1@nasa.gov.

Bryan Green is a senior software engineer at Grid Net. He was a member of the Visualization Group at the NASA Ames Research Center's Advanced Supercomputing facility, where he contributed to key CAMVis system technologies, including developing software for concurrent visualization. His current research interests include developing application firmware for smart-meter devices. Green has a BA in computer science from the University of California, Santa Cruz. Contact him at bryan@grid-net.com.



Evolving Critical Systems
 Edited by Mike Hinchey, Lorcan Coyle,
 Bashar Nuseibeh, and José Luis Fiadeiro

Computer
Essential Articles on Software Engineering

IEEE Computer Society
 CSPress

Easily browse this EssentialSet using buttons or bookmarks.

CONTENTS NEXT

NEW from  **CSPress**

EVOLVING CRITICAL SYSTEMS

Edited by Mike Hinchey, Lorcan Coyle, Bashar Nuseibeh, and José Luis Fiadeiro

This collection of articles from *Computer* magazine looks at critical software that can successfully evolve without loss of quality—while retaining its security, performance, reliability, and predictability.

PDF edition • \$29 list / \$19 members • 53 pp.

Order Online:
COMPUTER.ORG/STORE