

Recent Development of the NASA CAMVis for Tropical Cyclone Studies

Bo-Wen Shen^{1,2}(bo-wen.shen-1@nasa.gov), Wei-Kuo Tao², and Bron Nelson³

¹University of Maryland, College Park; ²NASA Goddard Space Flight Center; ³NASA Ames Research Center

Abstract

One of the current challenges in tropical cyclone (TC)¹ research is how to improve our understanding of TC inter-annual variability and the impact of climate change on TCs. Paired with the substantial computing power of the NASA Columbia and Pleiades supercomputers, the newly-developed Coupled Global Multiscale Modeling and Concurrent Visualization System (CAMVis; Shen et al., 2011) shows potential for such studies. The CAMVis consists of the NASA state-of-the-art multi-scale modeling framework (MMF; Tao et al., 2008, 2009), a high-resolution general circulation model (fvGCM; Shen et al., 2006), the Goddard Cumulus Ensemble model (GCE; Tao et al., 1993) and concurrent visualization (CV) systems (Ellsworth et al., 2006; Green et al., 2010). With the goal of improving hurricane climate simulations, a recent development within CAMVis is discussed that improves the scalability of the MMF making it feasible to perform long-term simulations and hence improve scientific visualizations of simulated hurricanes (e.g., Katrina) and consequently yield more insight into the understanding of hurricane transient dynamics. A meta grid system is introduced wherein thousands of copies of the GCE are integrated into a meta-global GCE. A revised parallelism is implemented with a 2D domain decomposition in this grid-point space. The mgGCE and fvGCM are then coupled into the new version 2 of the MMF. The revised parallel implementation in MMF v2 shows very promising scalability, giving a nearly linear speedup as the number of CPUs is increased from 30 to 364. In addition, a 3D streamline package is developed to provide insight into the multiscale interactions between a (mesoscale) hurricane and the large-scale environmental flow. 3D visualizations of hurricane Katrina show that the interaction of Katrina's outflow with an approaching upper-level jet stream may have been an important process that led to Katrina's intensification before it made landfall. The visualizations of Katrina and other TCs illustrate how CAMVis can help address the scenarios of "extreme event warning" to achieve the goal of "discovering predictive relationships between meteorological and climatological events and less obvious precursor conditions from massive data sets." The improved CAMVis makes it more feasible to study TC climate and has the potential to support the Decadal Survey Missions (NRC, 2007). Future work on model improvement will be discussed at the end.

¹ Depending on their location, TCs are referred to by other names, such as hurricane (in the Atlantic region), typhoon (in the West Pacific region), tropical storm, cyclonic storm, and tropical depression.

1. Introduction

Studies in TC inter-annual variability and the impact of climate change (e.g., global warming) on TCs have received increasing attention (e.g., Kerr, 2006), particularly due to the fact that 2004 and 2005 were the most active hurricane seasons in the Atlantic while 2006 was not as active as predicted. In addition, during the past ten years, statistics have shown that hurricane is the second deadliest weather event in the USA (Figure 1). Therefore, there is an urgent need of improving short-term and long-term hurricane forecasts. Thanks to recent advancements in global numerical models and supercomputer technology, these topics can be addressed better than ever before.

Earth (atmospheric) modeling activities have been conventionally divided into three major categories based on scale separations: synoptic-scale (large-scale), meso-scale (medium-scale), and cloud (micro)-scale. Historically, partly due to limited access to computing resources, TC climate has been studied mainly with general circulation models (GCMs) (Bengtsson et al., 2007) and partly with regional mesoscale models (MMs). The former have the advantage of simulating global large-scale flow, while the latter make it possible to simulate realistic TC intensity and structure with fine grid spacing. However for TC climate studies, the resolutions that were previously used in GCMs and MMs were still too coarse to resolve small-scale convective motion, and therefore "cumulus parameterizations" (CPs) were required to emulate the effects of unresolved subgrid-scale motion. Because the development of CPs has been slow, their performance is a major limiting factor in TC simulations.

Cloud-resolving models (CRMs) have been extensively developed to accurately represent non-hydrostatic cloud-scale convection and its interaction with environmental flows, aimed at improving TC prediction and advancing the development of CPs. Recently, an innovative approach that applies a massive number of CRMs in a global environment has been proposed and used to overcome the CP deadlock in GCMs (Randall et al., 2003; Tao et al., 2008, 2009). This approach is called the multiscale modeling framework (MMF) or super-parameterization, wherein a CRM is used to replace the conventional CP at each grid point of a GCM. Therefore, the MMF has the combined advantages of the global coverage of a GCM and the sophisticated microphysical processes of a CRM and can be viewed as an alternative to a global CRM. Currently, two MMFs with different GCMs and CRMs have been successfully developed at CMAP and NASA Goddard Space Flight Center (GSFC), and both have produced encouraging results in terms of a positive impact on simulations of large-scale flows via the feedback of explicitly resolved convection by CRMs. Among them is the improved simulation of the Madden-Julian Oscillation (MJO; Madden and Julian, 1970),

which could modulate TC activities. However, this approach poses great challenges in computing resources, data storages and data analyses for multi-decadal simulations of global TC activities. These challenges include but not limited to (1) running nearly 10,000 instantiations of the CRM with a great parallel performance; (2) efficiently archiving massive volume of data; (3) effectively conducting analyses in order to discover the predictive relationship between the meteorological (short-term) and climatological (long-term) events; (4) improving MMF’s scalability in order to increase GCM’s resolution for capturing realistic TC structure. To address the related issues, a revised model coupling approach is proposed to improve the Goddard MMF’s parallel scalability for TC study and a quasi 3D streamline package is developed to improve the representation of the sophisticated multiscale interactions during the hurricane’s formation, intensification and movement from massive volume of model outputs.




Figure 1: Statistics of fatalities caused by extreme weather events during the past 10 (in red) and 30 (in light blue) years. The 10-year statistics show that hurricane is the second deadliest event.

In section 2 of this article, we introduce the supercomputing, concurrent visualization (CV) and global modeling technologies at NASA. We then discuss in section 3 how a revised parallelism is implemented in the MMF of the CAMVis system. In section 4, the development of a quasi 3D streamline package is discussed. In section 5, we illustrate that 3D visualizations of Hurricane Katrina and a pair of twin TC in May 2002 can provide insightful understanding of multiple physical processes and their multi-scale interactions, with the aim of discovering the predictive relationship between TC activities and environmental flows. We conclude with a summary and discussion of future plans in section 6.

2. The NASA Supercomputers, CV, Goddard MMF

2.1 NASA Supercomputers

In late 2004, the Columbia supercomputer (Biswas et al., 2007) came into operation at NASA Ames Research Center. It consists of twenty 512-CPU nodes, giving it 10,240 CPUs and 20

terabytes (TB) of memory. Columbia achieved a performance of 51.9 Tflop/s (trillion floating point operations per second) with the LINPACK (Linear Algebra PACKage) benchmark. These large-scale computing capabilities enable complex problems to be resolved with large-scale modeling systems (e.g., Tao et al. 2008; Shen et al. 2006). In late 2008, the Pleiades supercomputer, an SGI Altix ICE system with a peak performance of 609 Tflop/s, was built as one of the most powerful general-purpose supercomputers. Recently, Pleiades has been upgraded to have 81,920 cores with a peak performance of 772.7 Tflops, 127 TB memory, and 3.1 Peta byte disk space. This newly built system, which provides more than 10 times the computing power of Columbia, is expected to speed up scientific discovery at an unprecedented pace. However, three different CPU architectures on Pleiades, which are summarized in Table 1, post challenges in optimizing the model’s performance.

Table 1: Major specifications of processor architects on Pleiades and Columbia, which might have impact on the parallel performance and scalability of the CAMVis and its components.

	cpu	Cores per processor	Cores per node	Clock speed	Memory per core	cache
Columbia	Itanium-2	1*	4 (512)	1.5 GHz	2GB	6MB
Pleiades	Harpertown (Xeon E5472)	4	8	3.0 GHz	1GB	6MB per pair of cores; no L3 cache
Pleiades	Nehalem (Xeon X5570)	4	8	2.93 GHz	3GB	8MB L3 cache shared by the four cores in one processor
Pleiades	Westmere (Xeon X5670)	6	12	2.93 GHz	2GB	12MB L3 cache shared by the six cores

2.2 Concurrent Visualization (CV) System

Large-scale Earth modeling systems enabled by the modern supercomputers poses great challenges to stage, handle, and manage these model outputs and compare them with satellite data. To overcome these challenges, concurrent visualization (CV) technology (Ellsworth et al., 2006; Green et al., 2010) has been developed at NASA/ARC. 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 has the benefit of providing much higher temporal resolution than is possible with traditional post-processing, enabling every timestep 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 may be useful for application monitoring or steering. This can help detect serious job and avoid wasting system resources.

CV technology was first developed and integrated into the high-resolution fvGCM on the original hyperwall system (49 screens). Recently, a new improved CV system (Version 2; Green et al., 2010) has been deployed on NASA’s 128-screen hyperwall-2, which is capable of rendering one-quarter-billion pixel graphics. CV 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. The 128-screen hyperwall-2, which is fully integrated into the NASA supercomputing environment, has a modern graphics card,

InfiniBand interconnects, 1,024-CPU cores and 475 terabytes of fast disk. These unique features provide an excellent environment for parallel feature extraction and extract storage. In addition, hyperwall-2's high-speed interconnect makes fully 3D concurrent visualization possible.

To efficiently exchange data between the computing and visualization nodes, we have implemented the M-on-N configuration for the CV pipeline. The M-on-N configuration allows different domain decomposition within computing and visualization nodes. More details on the new features of the CV version 2.0 can be found in Green et al. (2010) and Shen et al. (2011).

2.3 Multiscale Modeling Framework (MMF)

The Goddard MMF is based on the NASA Goddard finite-volume GCM (fvGCM) and the Goddard Cumulus Ensemble model (GCE). While the high-resolution fvGCM has shown remarkable capabilities in simulating large-scale flows and thus hurricane tracks (Atlas et al., 2005; Shen et al. 2006a-b, 2010a-b, 2011), the GCE is well known for its superior performance in representing small cloud-scale motions and has been used to produce more than 90 refereed journal papers (Tao et al., 1993, 2003). In the MMF, the fvGCM is running at a coarse ($2^\circ \times 2.5^\circ$) resolution, and one copy of GCE is running within each of the fvGCM grids (Figure 2). This resolution has grid points of (144, 91) in the (x, y) directions. This gives a total of 144x90 (12,960) horizontal cells. Thus, 12,960 GCEs are “embedded” in the fvGCM to allow explicit simulation of cloud processes in a global environment. Currently, only averaged thermodynamic fields such as temperature and water vapor in the GCE are fed back to the fvGCM, which is called “thermodynamic feedback”. The time step for the individual 2D GCE is ten seconds, and the fvGCM-GCE coupling interval is one hour at this resolution. Under this configuration, 95% or more of the total wall-time for running the MMF is spent on the GCEs. Thus, wall-time could be significantly reduced by efficiently distributing the large number of GCEs over a massive number of processors on a supercomputer.




Figure 2: The grid systems of the fvGCM (left) and the embedded GCE (right). fvGCM has a 3D grid in the (x, y, z) direction in association with (longitude, latitude, and altitude), while the GCE has a 2D grid in the (x, z) direction. In this configuration, a GCE with 64x32 grid points in the (x, z) direction is embedded in each of fvGCM grids and is used to generate the statistics of cloud activities averaged over the entire GCE horizontal extent.

Over the past few years, an SPMD (single program multiple data) parallelism has been separately implemented in both the fvGCM and GCE with good parallel efficiency (Putman et al. 2005; Juang et al. 2007). However, the key for improving the overall performance is to increase the copies of the GCEs to be run in parallel, which will be discussed in details later. We will first introduce the GCE and fvGCM and then discuss a revised strategy for coupling the fvGCM and massive copies of the GCE with improved scalability.

2.3.1 The Goddard Cumulus Ensemble (GCE) model

Over the last two decades, the Goddard Cumulus Ensemble model (GCE) has been developed in the mesoscale dynamics and modeling group, led by Dr. W.-K. Tao, at NASA Goddard Space Flight Center. The GCE has been well tested and continuously improved, and its main features were described in detail in Tao et al., (1993, 2003). Typical model runtime configurations for MMF runs are: (a) 64 grid points in the x direction with a grid spacing of 4 km; (b) 32 vertical stretched levels; (c) cyclic lateral boundary conditions; and (d) a time step of 10 seconds.

The GCE has been implemented with a 2D domain decomposition using MPI-1 (Message Passing Interface version 1) to take advantage of recent advances in supercomputing power (Juang et al., 2007). To minimize the changes in the GCE, implementation was done with a separate layer added for data communication, which preserves all of the original array indices. Therefore, not only code readability for existing modelers/users but also code portability for computational researchers is maintained. In addition to “efficiency” enhancement, tremendous efforts were made to ensure reproducibility in simulations with different CPU layouts. Without this, it would be difficult for model developers to test the model with new changes and to compare long-term simulations generated with different numbers of CPUs. However, the current MMF is not yet able to take advantage of the fine-grain parallelism inside the GCE, which can be used after a revised coarse-grain parallelism is implemented and well tested in the new MMF.

2.3.2 The finite-volume General Circulation Model (fvGCM)

Resulting from a development effort of more than ten years, the finite-volume General Circulation Model (fvGCM) is a unified numerical weather prediction (NWP) and climate model that can run on daily, monthly, decadal, or century time-scales. It has the following major components: (1) finite-volume dynamics (Lin et al., 2004), (2) physics packages from the NCAR Community Climate Model Version 3 (CCM3), and (3) the NCAR Community Land Model Version 2 (CLM2). The model was originally designed for climate studies at a coarse resolution of about 2×2.5 degree in the 1990s, and its resolution was increased to 1 degree in 2000 and 1/2 degree in 2002 for NWP. Since 2005, the ultra-high resolution (e.g., 1/8 and 1/12 degree) fvGCM has been deployed on the Columbia supercomputer, showing remarkable TC forecasts.

The parallelization of the fvGCM was carefully designed to achieve efficiency, parallel scalability, flexibility, and portability. Its implementation had a distributed- and shared-memory two-level parallelism, including a coarse grained parallelism with MPI² (MPI-1, MPI-2, MLP, or SHMEM) and fine grained parallelism with OpenMP (Putman et al. 2005). The model’s dynamics, which require a lot of inter-processor communication, have 1D MPI/MLP/SHMEM domain decomposition in the y direction and OpenMP multithreading in the z direction. One of the prominent features in the implementation is to allow multi-threaded data communication. The physical part was parallelized with the 1D domain decomposition in the y direction inherited from the dynamics part and further enhanced with an OpenMP loop-level parallelism in the decomposed latitudes. CLM2 was also implemented with both MPI and OpenMP parallelism, allowing its grid cells to be distributed among processors. Between dynamical grid cells and land patches, a data mapping (or redistribution) is required.

Though the fvGCM has MPI for coarse grain parallelism and OpenMP for finer grain parallelism, the latter is not applicable for the current MMF runs. In addition, as the MPI parallelism is applied to a one-dimensional decomposition over latitude, and each MPI task needs at least three grid points in latitude for parallel efficiency reasons, the MPI parallelism is very limited for small computational grids (Figure 3). For the 2x2.5 degree grids, 30 MPI tasks are the maximum.

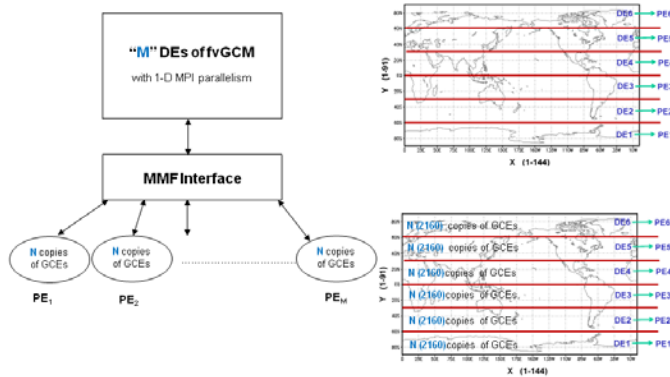


Figure 3: Parallelism in the fvGCM and the fvMMF v1.0. The fvGCM has the MPI parallel implementation with an 1-D domain decomposition along with the y direction. The approach of embedding one copy of GCE into each of the fvGCM’s grids inherits the parallelism and thus limited scalability of the fvGCM. DE stands for domain element, and PE for processing element. The right panels show the distribution of DEs by using M=6 PEs, which gives 2,160 copies of GCEs running in a PE.

3. Revised parallelism in the MMF

In the MMF version 1, each of 12,960 GCEs is embedded into a grid cell of the fvGCM. Though this implementation is

² To simplify discussion in this article, the term “MPI” used along with the fvGCM will be referred to as any one of MPI-1/MPI-2/MLP/SHMEM communication paradigms.

straightforward by adopting the parallel framework of the fvGCM, it inherits both the advantages and disadvantages of the fvGCM framework. One of the disadvantages is that this approach limited the parallel scalability of the MMF for small domain grids, which requires many copies of GCEs to run within one processing element (PE). For example, 2,160 copies of GCE are running in series in one PE when only 6 PEs are used, as shown in Figure 3. Even though the maximum PEs (30) are used, each PE still needs to run 432 copies of GCEs! Therefore, how to make more copies of GCEs to run in parallel with more PEs is the key to reduce the wall-clock time of running MMF, which could make it more feasible to perform long-term climate simulations. It could also make it practical to increase the resolution of the fvGCM and/or extending the GCE’s dimension from 2D to 3D. To achieve these goals, a different strategic approach is proposed to couple the fvGCM and GCEs.

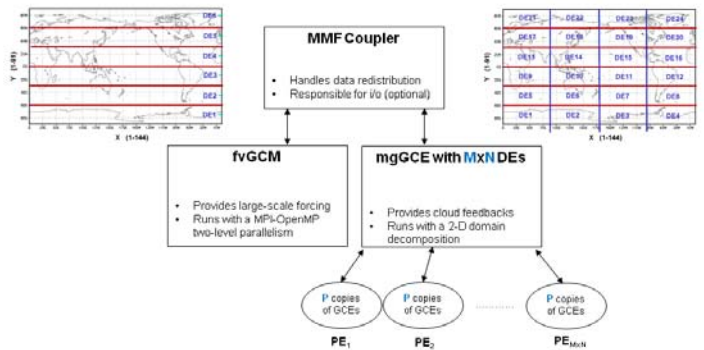


Figure 4: Revised parallelism in the new version of the fvMMF. The original parallel implementation in the fvGCM remains unchanged. The massive copies of GCEs are “integrated” into a super-component called meta-global GCE (mgGCE) in which a 2D domain decomposition is applied. As domain decompositions in the fvGCM and mgGCE are different, a coupler is developed to handle data redistribution among these two components. This revised parallelism improves the scalability of fvMMF by allowing more copies of GCEs running in parallel, and thus reduces wall-time significantly because all of the GCEs cost 95% or higher of walltime for a MMF run.

From a computational perspective, the concept of “embedded GCEs” restricts the view on the data parallelism of the fvGCM. Because a periodic lateral boundary condition is used in a GCE, execution of each embedded GCE is independent of the other GCEs during a timestep of the fvGCM model, and thus it can be run in a separate MPI task. Accordingly, we propose a new coupling approach to improve the MMF’s parallel scalability. This approach integrates all copies of the GCEs into an individual component, reducing the dependence on the framework of the fvGCM. This super-component with 12,960 GCEs could be viewed as a meta global GCE (mgGCE) in a meta grid-point system, which includes 12,960 grid points. This grid system, which is not tied to any specific grid system, is assumed to be the same as the latitude-longitude grid structure in the fvGCM for convenience. With this concept in mind, each of the two individual components (the fvGCM and mgGCE) in the MMF could have its own scaling properties (Figure 4). Since most of

wall-time was spent on the multiple instantiations of the GCE, we could substantially reduce the wall-time by deploying a highly scalable mgGCE and coupling the mgGCE and the fvGCM using an MPMD (multiple programs multiple data) parallelism. For example, the 2D domain decomposition in the mgGCE enables more PEs to be used for running more copies of GCEs in parallel.

Thus, the parallel scalability and performance of the MMF will depend mainly on the parallel scalability and performance of the mgGCE and the coupler, which is the interface for data regridding between the fvGCM and mgGCE. Because cyclic lateral boundary conditions are used in each GCE, the mgGCE has no ghost region in the meta grid system and can be “embarrassing” scalable with a 2D domain decomposition, which can in turn greatly help improve MMF’s scalability. A data parallelism in the mgGCE is naturally a task parallelism, namely distributing 12,960 GCEs among processors. This is a coarse-grained parallelism as compared to the parallelism inside a GCE. Under this current definition, a grid inside each GCE which corresponds to one meta grid becomes a *child grid* (or sub-grid) with respect to the parent (meta) grid (Figure 2b). Since an individual GCE can still be executed with its native 2D MPI implementation in the child grid-point space, this second level of parallelism (fine-grained parallelism) can greatly expand the number of CPUs. Potentially, the coupled MMF along with the mgGCE could be scaled at a multiple of 12,960 CPUs. To achieve all of the aforementioned functionalities, we need to develop a scalable and flexible coupler and a scalable parallel I/O module. The coupler will be designed carefully in order: (1) to minimize the changes in the GCE and permit it as a stand-alone application with a single component of either fvGCM or the mgGCE; (2) to seamlessly couple the mgGCE and fvGCM to allow for a different CPU layout in each of these components; and (3) to allow the mgGCE to be executed in a global, channel, or regional environment with a suitable configuration (which is called a the cloud-mask file). A scalable (parallel) I/O module with a different CPU layout that could be different from that in the either fvGCM or the mgGCE will be implemented in the meta grid-point space, as it is impractical to have each instantiation of GCEs to do its I/O.

At this time, a prototype MMF including the mgGCE, fvGCM and coupler has been successfully implemented. The technical approaches are briefly summarized 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 by 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; (4) data gathering in the mgGCE is done along the x direction and then the y direction; and (5) synchronization is implemented with the atomic *__sync_add_and_fetch* function call on the Columbia supercomputer. While steps (1), (2), and (5) were previously used in MLP (multiple level parallelism) (Taft, 2001), this methodology is now extended to the multi-component system.

Figure 5 shows preliminary benchmarks with very promising parallel scalability up to 364 CPUs. Here the speedup is

determined by T_{30}/T , where T is the wall-time to perform a 5-day forecast with the MMF and T_{30} the time spent using 30 CPUs. The run with 30 CPUs was chosen as a baseline simply because this configuration was previously used for production runs (Tao et al., 2009). A speedup of (3.93, 7.28, and 12.43) is obtained by increasing the number of CPUs from 30 to 91, 182, and 364 CPUs, respectively. As the baseline has load imbalance and excessive memory usage in the master process, it is not too surprising to obtain a super-linear speedup. Further analysis of the MMF’s throughput indicates that it takes about 164 minutes to finish a 5-day forecast using 364 CPUs, which meets the requirement for performing realtime numerical weather prediction. A yearly simulation would only take 8 days to run with 364 CPUs as opposed to 96 days with 30 CPUs. This makes it far more feasible for studying TC climate. The enhanced coupled model has been used to perform two-year production runs.




Figure 5: Parallel scalability of the Goddard MMF with a revised parallel implementation on the NASA Columbia supercomputer. This figure shows that a linear speedup is obtained as the number of CPUs increases from 30 to 364. The original MMF could use only 30 CPUs. Further improvement is being conducted.

4. Recent Improvement of Concurrent Visualization System

4.1 Concurrent Visualization to Web (CV2Web)

In order to maximize the results from a single simulation run, multiple products are usually generated, representing various fields and regions of interest as well as numerous feature-extraction and visualization techniques. When time-stamped outputs arrive from the computing nodes, each visualization node sequentially computed all of the requested visualizations, producing one image per visualization request. As part of the CV pipeline, the resulting animations are streamed, as they are being generated, to the remote displays at the facilities of the investigators. Though these features provide unprecedented opportunities for researchers to conduct research activities, the dependence of supercomputer-scale visualization facilities seem to become an issue for many end users. Recently, in order to support more end users such as managers, concurrent visualization to Web (CV2Web, see details at <http://www.nas.nasa.gov/>) has been implemented to provide a

simple means for accessing real-time visualization results produced by the CAMVis. This new capability is built on a hyperwall concurrent visualization pipeline, enabling real-time web access to concurrent visualization products. The CV2Web leverages the InfiniBand fabric between Pleiades and hyperwall-2 to give users with instrumented codes the ability to see multiple visualizations of their simulation in progress via a web browser, from any location. While the current version supports real-time image updates for each visualization, subsequent versions will support animation or video.

4.2 Streamline Package (StrPack) in the CV

In order to seamlessly visualize all different fields and their predictive relationship from model simulations with the CAMVis, we have integrated and developed different visualization packages. One of them which can help visualize the sophisticated multiscale interaction of hurricane dynamics is the streamline package (StrPack). The basic concept of calculating the streamlines with the StrPack (Figure 6) is as follows: (1) given a (u,v) vector field, a streamline is produced by starting from a "seed" point, and (2) then tracing the path that a hypothetical particle would take through that field. Thus, the streamline is parallel to the vector field at each point. Seed points are typically chosen by selecting some subset of the (x, y) points that give rise to the (u, v) field. One might for instance select every 5th point in the x direction, and every 3rd point in the y direction. Selection of seed points is discretionary: pick enough that the resultant image has sufficient detail, but not so many that the image is cluttered. In contrast, when the granularity of the (u,v) field from a specific (small) subdomain is coarser than the number of pixels available in the image, the values are interpolated to produce smoother lines. Streamlines are an "instantaneous" property of the vector field, highlighting what the field looks like at that moment.




Figure 6: Methodology of making a streamline. (a) In a 2-D vector field (shown in blue), the direction of the arrow represents the direction of the vector at each point, and the length of the arrow represents the magnitude. (b) Select a "seed" point (indicated by the small square with a red border) to begin the streamline. (c) Trace the path a particle starting from the seed point would follow as it was pushed by the vector field. (d)-(e) repeat the step (c).

5. Scientific Applications with the CAMVis

In this section, we discuss scientific applications with the newly-developed CAMVis (Figure 7), which includes the newly-developed mgGCE with a revised parallelism and the newly-

developed CV 2.0 with a capability for M-on-N parallel data transfer. In 2007, the National Research Council (NRC) Decadal Survey entitled: "Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond" was completed at the request of NASA and other government agencies. The Decadal Survey recommends that "The U.S. 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." To plan a technology roadmap in support of the missions recommended by the Decadal Survey, the NASA Earth Science and Technology Office (ESTO) held an AIST road-mapping workshop in Cocoa Beach, FL, February 8-11 2010. During the workshop, there were three group meetings for the discussion of the scenarios in the Decadal Survey. These groups are Sensor System Support, Advanced Data Processing (ADP) and Data Services Management. The ADP group identified "Extreme Event Warning" and "Climate Prediction" as two of the top priority scenarios. The first focuses on "discovering predictive relationships between meteorological and climatological events and less obvious precursor conditions from massive data sets. The second is to seek "robust estimates of primary climate forcings for improved climate forecasts, including local predictions of the effects of climate change." In the following, we discuss the functionalities of the CAMVis in support of the "Extreme event warning" with the aim of improving our confidence in the the model's performance for climate prediction.

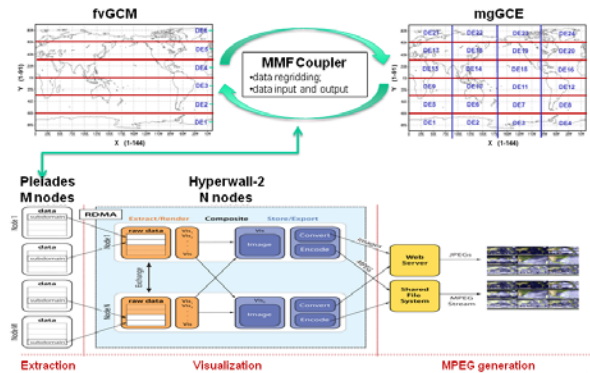


Figure 7: A diagram of the CAMVis 1.0, which consists of the MMF (=fvGCM + mgGCE, top) running with M nodes on Pleiades and CV 2.0 (bottom) running with N nodes on Hyperwall-2. A MMF coupler is developed to enable different parallel implementations in the fvGCM and mgGCE and perform regridding between them. A M-on-N parallel data transfer is implemented between the MMF and the CV.

5.1 Visualizations of multiscale interactions for Katrina

During the past 20 years, statistics of hurricane forecasts by the National Hurricane Center show that hurricane track forecasts have been improving steadily, but hurricane intensity prediction has lagged behind. It is known that improving hurricane intensity with a global model would require fine grid spacing and realistic model physics to simulate the fine structure of a hurricane and the interaction with its environmental flows. It was very challenging

with traditional global models. The recent advances in supercomputing and global model technologies at NASA have provided an unprecedented opportunity of improving our understanding of the hierarchical multiscale interactions of Hurricanes and tropical cyclones and thus of improving their predictions.

It is known that accurate simulations of hurricane structure are crucial for intensity predictions. The key features of the hurricane structure (Figure 8) include (1) a hurricane eye: a region in the center of a hurricane (tropical cyclone) where the winds are light and skies are clear to partly cloudy; (2) a hurricane eyewall: a wall of dense thunderstorms that surrounds the eye of a hurricane; (3) low-level inflow with counter clockwise circulation; (4) upper-level outflow with clockwise circulation; and (5) an elevated warm core (not shown) where the cyclone's temperature is warmer at its center than at its periphery. In addition to the hurricane structure, accurate representation of the scale interaction between a hurricane and its environmental flows is another major factor that impacts on hurricane's intensity. For example, Figure 9 suggests that the interaction of the upper-level jet stream and the outflow of a hurricane with clockwise circulation may lead to the intensification of the hurricane.




Figure 8: Major characteristics of a hurricane, including (1) an eye: a region in the center of a hurricane (tropical cyclone) where the winds are light and skies are clear to partly cloudy; (2) an eyewall: a wall of dense thunderstorms that surrounds the eye of a hurricane; (3) low-level inflow with counter clockwise circulation; (4) upper-level outflow with clockwise circulation; (5) an elevated warm core (not shown) where the cyclone's temperature is warmer at its center than at its periphery. (courtesy of the COMET program at NCAR)




Figure 9: A schematic diagram of the interaction between the upper-level hurricane outflow (with clockwise circulation) and a jet stream, which could lead to the intensification of the hurricane. (courtesy of the COMET program at NCAR)

Enabled by the NASA supercomputing technologies, we first deployed the high-resolution fvGCM at a resolution of 1/8 degree on the Columbia supercomputer and obtained realistic simulation of Katrina's movement, intensity and near-eye wind distribution from a 5-day run (e.g., Shen et al., 2006a). With the support from the AIST program, we now successfully deployed the 1/8 degree fvGCM on the Pleiades supercomputer which is a distributed-memory platform. We obtained comparable results (e.g., 96h simulation of the surface near-eye winds) on the Columbia and Pleiades supercomputers, with the ending date of 1200 UTC August 29, 2005 prior to Katrina's landfall (not shown).

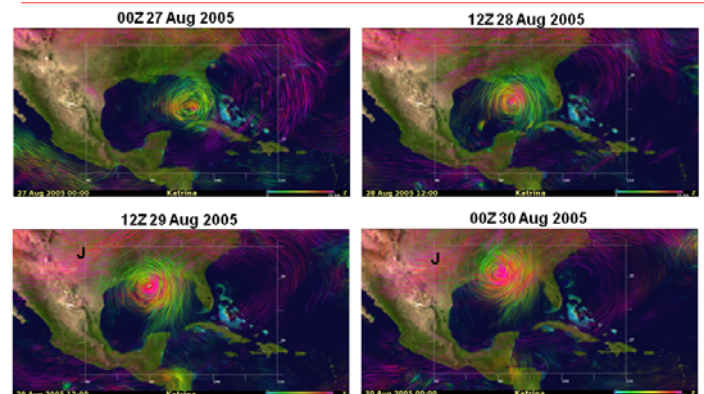


Figure 10: Streamline visualization of multiscale interaction between the outflows of Hurricane Katrina and an approaching upper-level jet stream from a 5-day 1/8 degree run, which is associated with the intensification of Katrina before landfall. Low level streamlines are in blue, while upper level streamlines in pink and red. Katrina's intensification is indicated by the appearance of dense red streamlines at the upper levels (panels c and d), which are in association with strong vertical motion. The symbol 'J' indicates the jet stream.

With the aforementioned encouraging results for Katrina's structure, it becomes feasible to examine the role of the interaction between Katrina's outflow and its environmental flows on its intensification. Figure 10 shows the evolution of a simulated Katrina from a 5-day run. All of the frames are derived from a high-resolution visualization, which was discussed during the annual project review. Low-level winds are in blue and upper-level winds in red. As time progresses, stronger upper-level flows developed in association of the initial intensification of Katrina (e.g., Figure 10b). Later on, Katrina experienced rapid intensification as its upper-level flows interacted with an approaching upper-level jet stream (Figure 10c,d). It should be noted that quasi 3D streamlines seem to be very effective in representing the scale interaction and linking it to the Katrina's intensification. *However, as the streamline packages were just finished, more rigorous verification on the accuracy of scientific representation are still needed before we can claim that this multiscale interaction may contribute to the*

intensification of Katrina. Related tasks are being planned and will be conducted in the third year of the AIST project.

5.2 Visualizations of formation for twin tropical cyclones

Improving understanding of hurricane formation is one of the most challenging hurricane research activities. This is also crucial for improving understanding of hurricane climate. During the past two years with the support by the AIST project, the performance of the high-resolution fvGCM in simulating the formation of hurricane and tropical cyclones has been verified against the global analyses and NASA Satellite data such as QuikSCAT sea winds and TRMM precipitations. These tropical cyclones include the severe cyclonic storm Nargis (2008) by Shen et al. (2010a) and Hurricane Helene (2006) by Shen et al. (2010b), and tropical cyclones in May 2002 by Shen et al. (2011b, to be submitted). As discussed earlier that the CAMVis is capable of examining the transient dynamics during the model integration, we illustrate this with the twin TC. While the high temporal resolution visualization was presented during the annual project review, individual frames are used here for the written report.

Figure 11 with 4 panels derived from the visualization of twin TC in May 2002 shows their formation. In each panel, the Northern and Southern Hemispheres are in the left and right **side of the panel**, respectively. Labels ‘E’ and ‘W’ indicate easterly and westerly winds, respectively. Panel (a) shows the formation of the TC Kesiny, the southern part of the twin TC. As time proceeds, the interaction between the westerly and easterly winds in panels (b) and (c) may lead to the formation of TC01A, the northern part of the twin TC in panel (d). The white arrow in each of the zoomed-in panels is referred to as the spinning axis, and roughly indicates the location of vortex centers at different heights. Through a high temporal resolution visualization, it can be shown that the spinning axis of a mature TC points vertically (in the z direction), while the direction of the spinning axis changes with time during the formation stage. The latter suggests that the vortex centers at different heights are not coherent, which suggests it is challenging to improve the initialization of a weak vortex because its vertical coherence is absent.

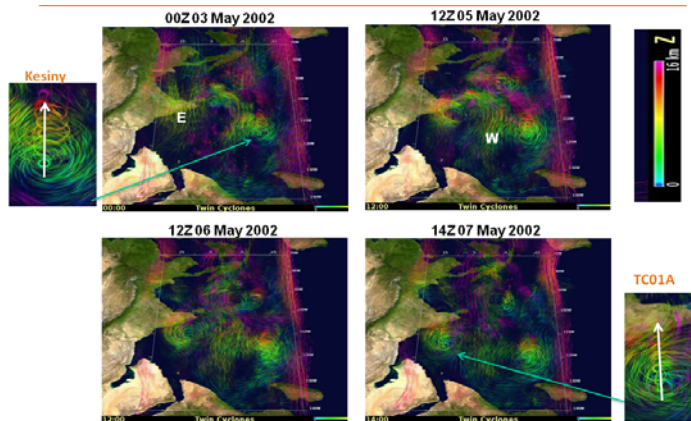


Figure 11: Visualization of the formation for the twin TC in May 2002. In the view of each panel, the Northern and Southern Hemispheres (NH and SH) are in the left and right side of the

panel, respectively. Labels ‘E’ and ‘W’ indicate easterly and westerly winds, respectively. Panel (a) shows the formation of the TC Kesiny in the SH. The interaction between the westerly and easterly winds in panels (b) and (c) may lead to the formation of TC01A (d), the counter part of the twin TC in the NH. The white arrow (named as the spinning axis) in each of the zoomed-in panels indicates the location of the vortex centers at different heights. Through a high temporal resolution visualization, it can be shown that the spinning axis of a mature TC points vertically (in the z direction), while the direction of the spinning axis changes with time during the formation stage.

6. Concluding Remarks

Improving our understanding of TC inter-annual variability and the impact of climate change (e.g., doubling CO₂ and/or global warming) on TCs brings both scientific and computational challenges to researchers. As TC dynamics involves multiscale interactions among synoptic-scale flows, mesoscale vortices, and small-scale cloud motions, an ideal numerical model suitable for TC studies should demonstrate its capabilities in simulating these interactions. The newly-developed Coupled Advanced Multiscale modeling and concurrent Visualization systems (CAMVis; Shen et al., 2011) on the NASA Columbia (Biswas et al., 2007) and Pleiades supercomputers show promise in pursuing the related studies. The CAMVis consists of the state-of-the-art multiscale modeling framework (MMF; Tao et al., 2008,2009), fvGCM (Shen et al. 2006a-b; Shen et al. 2010a-b) and 2D GCEs (Tao et al. 1993; 2003), and concurrent visualization systems (Ellesworth et al., 2006; Green et al., 2010). This article focuses on the recent development and improvement of the CAMVis during the execution of the AIST project for the second year, which include (1) the development of the meta-grid GCM (mgGCE); (2) a revised parallel implementation to improve the MMF’s performance and parallel scalability; (3) deployment of the 1/8 degree fvGCM on the Pleiades supercomputer; (4) improvement of the parallel “M-on-N” data transfer model in the CV system version 2.0, which enables parallel data transfer between “M” computing nodes on Pleiades and “N” visualization nodes on Hyperwall-2; (5) development of the CV to Web (CV2Web) that is a new capability of the hyperwall-based concurrent visualization pipeline, enabling real-time web access to concurrent visualization products; (6) development of the quasi 3D streamline packages (StrPack) to provide insightful understanding of the hurricane’s multiscale interactions and transient dynamics; and (7) development of data modules to fuse NASA satellite data such as QuikSCAT sea winds and TRMM precipitation for inter-comparisons with model simulations at comparable resolutions.

Our recent benchmarks show that the revised parallel implementation can improve the MMF’s scalability up to hundreds of CPUs without the need of major changes in the fvGCM and GCEs, making it more feasible to perform long-term climate simulations for hurricanes and tropical cyclones. Further performance and scalability improvements are being conducted. The ultra-high resolution fvGCM was first deployed on

Columbia, which was only one of a few global models running a resolution of 10km in 2005, to produce remarkable 5-day forecasts of Katrina's track and intensify. With the support by AIST, efforts have been put in place to deploy the 1/8 degree model on Pleiades, examine the impact of multi-core processors on the memory-intensive numerical model, improve our understanding of model's performance in hurricane predictions with 3D visualization and thus improve our confidence in the model's capability in long-term climate simulations. For example, applying the StrPack, we examine the multiscale interactions of hurricane Katrina with its upper-level jet stream, and illustrate the differences of 3D vortex dynamics between a weak tropical cyclone (such as the formation of twin TCs) and a strong TC (such as the Katrina at a mature stage). The related tasks improve our confidence in the model's performance in simulating TCs and thus could eventually provide justifications of performing hurricane climate simulations at decadal scales with the CAMVis.

The fundamental change in the MMF v2 that improves MMF's scalability is the development of the super-component mgGCE on a meta grid system that groups a large number of GCEs. This permits a component-based programming paradigm with which the fvGCM and mgGCE are coupled with their own parallelism. In the current version, the fvGCM has a parallelism with a 1D domain decomposition and the mgGCE with a 2D domain decomposition. A prototype coupler is then implemented in the MMF v2 for data redistribution between these two components. As 95% of the computing time for the MMF is spent on the mgGCE, the revised parallelism leads to a substantial performance. For example, a speedup of 12 is obtained by increasing the number of CPUs from 30 to 364 where 30 is the upper limit of CPU counts in the first version of MMF. The underlying communication paradigm for data redistribution in the MMF v2 is similar to the multiple-level parallelism (MLP, Taft 2001), which was previously used for parallelization in single-component models with tremendous benefits. The methodology is extended here to a multi-component modeling system, showing an alternative and easy way for coupling multiple components. Further improvements in the implementation include an adoption of a more portable communication paradigm (such as MPI-1 or MPI-2) and/or a sophisticated modeling framework. While the current implementation in process management, data communication/redistribution, and synchronization is solely done with Unix system calls, earlier experiences with the parallel implementation in the fvGCM have proven that this can be easily extended with an MPI-1 or MPI-2 implementation (e.g., Putman, Lin, and Shen, 2005). Further improvement in scalability and implementation with MPI communication protocols to scale the model up to thousands of cores/CPU are being planned and conducted.

This proof-of-concept approach lays the groundwork for a more sophisticated modeling framework and coupler to solve unprecedentedly complex problems with advanced computing power. For example, if computing resources are limited, a cloud-mask file can be used to specify limited regions where the GCEs should be running. A more sophisticated cloud-mask

implementation in the mgGCE will enable one to choose a variety of GCEs (2D vs. 3D) depending on geographic location. Thus, computational load balances can be manageable. It is well known that a latitude-longitude grid system has issues such as efficiency/performance and convergence problems near the poles. As the meta grid system in the mgGCM is no longer bound to the fvGCM's grid system, this meta-grid concept could help avoid the performance issues by implementing a quasi-uniform grid system (such as a cube grid or geodesic grid) into the mgGCE.

As a stand-alone model, the mgGCE can be also tested offline with large-scale forcing derived from model reanalysis [e.g., from the Global Forecast System (GFS) at the National Centers for Environmental Prediction (NCEP)] or from high-resolution model forecasts (e.g., from the fvGCM). To assure the implementation in the mgGCE to be correct, simulations with the mgGCE at a single meta point should be identical to those with a regular GCE. One of potential applications with the mgGCE is to investigate the short-term evolution of hurricane Katrina's (2005) precipitation by performing simulations driven by the NCEP GFS T382 (~35km) reanalysis data at a 6h time interval. Then, we can extend this approach by replacing the GFS reanalysis with 1/8° fvGCM forecasts at a smaller time interval (see more detailed information about these forecasts in Shen et al., 2006a). However, all of the aforementioned implementations and tasks are subject to a future study.

In this report, we discuss the development of the streamline package (StrPack) and its application to examining the scale interactions and transient dynamics of hurricane Katrina and twin TCs (2002). However, the current version of StrPack does not currently use the information about vertical wind velocity in the "z" dimension (altitude), mainly because the atmosphere is basically hydrostatic with a horizontal scale on the order of tens or hundreds of kilometers but a vertical extent up to 10-40 km. Consequently, the streamlines produced in the visualizations are not true 3D lines. Rather, each pressure level is treated independently: 2D streamlines are produced within that level, and then the levels are stacked to produce a quasi-3D image. We would like to do true 3D simulation and visualization, although it will probably take considerable work to improve the tools to accomplish this goal of greater realism. It will also take significant efforts to know exactly how much scientific understanding would be improved by doing so; the current approximation seems to provide good results in most cases.

In support of the Decadal Survey missions, the NASA Earth Science and Technology Office (ESTO) held an AIST road-mapping workshop. During the workshop, the Advanced Data Processing (ADP) identified "*Extreme Event Warning*" and "*Climate Prediction*" as two of the top priority scenarios. In the report, we discuss the functionalities of the CAMVis in simulating and visualizing the predictive relationship between a hurricane and its environmental flows such as an upper-level jet stream and tropical waves, illustrating the potential for addressing the scenario of "*Extreme event warnings*." Further improvement with scheduled tasks will be made to improve our confidence in the the model's performance for climate prediction.

Acknowledgements:

We are grateful to the following organizations for their support: the NASA Earth Science Technology Office, the Advanced Information Systems Technology Program, the NSF 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 ARC, and the NASA Center for Climate Simulation at GSFC.

References

- Atlas, R., O. Reale, B.-W. Shen, S.-J. Lin, J.-D. Chern, W. Putman, T. Lee, K.-S. Yeh, M. Bosilovich, and J. Radakovich, 2005: Hurricane forecasting with the high-resolution NASA finite volume general circulation model. *Geophys. Res. Lett.*, **32**, L03801, doi:10.1029/2004GL021513.
- Bengtsson, L., K., I. Hodges, and M. Esch, 2007: Tropical cyclones in a T159 resolution global climate model: comparison with observations and re-analyses. *Tellus A* 59 (4), 396-416 doi:10.1111/j.1600-0870.2007.00236.x
- Biswas, R., M.J. Aftosmis, C. Kiris, and B.-W. Shen, 2007: Petascale Computing: Impact on Future NASA Missions. *Petascale Computing: Architectures and Algorithms*, 29-46 (D. Bader, ed.), Chapman and Hall / CRC Press, Boca Raton, FL.
- Ellsworth D., B. Green, C. Henze, P. Moran, T. Sandstrom, 2006: Concurrent Visualization in a Production Supercomputing Environment. *IEEE Trans. on Visualization and Computer Graphics.*, **12**, 5, September/October 2006.
- Green, B., C. Henze, B.-W. Shen, 2010: Development of a scalable concurrent visualization approach for high temporal- and spatial-resolution models. AGU 2010 Western Pacific Geophysics Meeting, Taipei, Taiwan, June 22-25, 2010.
- Juang, J.-M., W.-K. Tao, X. Zeng, C.-L. Shie, S. Lang, and J. Simpson, 2007: Parallelization of NASA Goddard Cloud Ensemble Model for Massively Parallel Computing. *Terrestrial, Atmospheric and Oceanic Sciences*.
- Kerr, R., 2006: Sharpening Up Models for a Better View of the Atmosphere. *Science*, **313**, 1040.
- Lin, S.-J., 2004: A "vertically Lagrangian" finite-volume dynamical core for global models. *Mon. Wea. Rev.*, **132**, 2293-2307.
- Madden, R. A., and P. R. Julian, 1971: Detection of a 40-50 day oscillation in the zonal wind in the tropical Pacific. *J. Atmos. Sci.*, **28**, 702-708.
- National Research Council (NRC) Decadal Survey, 2007: *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond*. The National Academies Press, Washington, D.C. This decadal survey, which was completed at the request of NASA and other government agencies, recommends that: "The U.S. 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."
- Putman, W., S.-J. Lin, and B.-W. Shen, 2005: Cross-Platform Performance of a Portable Communication Module and the NASA Finite Volume General Circulation Model. *International Journal of High Performance Computing Applications*. **19**: 213-223.
- Randall, D., M. Khairoutdinov, A. Arakawa, W. Grabowski, 2003b: Breaking the Cloud Parameterization Deadlock. *Bull. Amer. Meteor. Soc.*, 1547-1564.
- Shen, B.-W., R. Atlas, J.-D. Chern, O. Reale, S.-J. Lin, T. Lee, and J. Chang 2006a: The 0.125 degree Finite Volume General Mesoscale Circulation Model: Preliminary simulations of mesoscale vortices. *Geophys. Res. Lett.*, **33**, L05801, doi:10.1029/2005GL024594.
- Shen, B.-W., R. Atlas, O. Reale, S.-J. Lin, J.-D. Chern, J. Chang, C. Henze, and J.-L. Li, 2006b: Hurricane Forecasts with a Global Mesoscale-Resolving Model: Preliminary Results with Hurricane Katrina (2005). *Geophys. Res. Lett.*, **33**, L13813, doi:10.1029/2006GL026143.
- Shen, B.-W., W.-K. Tao, W. K. Lau, R. Atlas, 2010a: Predicting Tropical Cyclogenesis with a Global Mesoscale Model: Hierarchical Multiscale Interactions During the Formation of Tropical Cyclone Nargis (2008). *J. Geophys. Res.*, **115**, D14102, doi:10.1029/2009JD013140.
- Shen, B.-W., W.-K. Tao, and M.-L. C. Wu, 2010b: [African Easterly Waves in 30-day High-resolution Global Simulations: A Case Study during the 2006 NAMMA Period.](#) *Geophys. Res. Lett.*, L18803, doi:10.1029/2010GL044355.
- Shen, B.-W., W.-K. Tao, and B. Green, 2011: Coupling NASA Advanced Multi-Scale Modeling and Concurrent Visualization Systems for Improving Predictions of Tropical High-Impact Weather (CAMVis), Computing in Science and Engineering (CiSE), 23 November 2010. IEEE computer Society Digital Library. IEEE Computer Society, <http://doi.ieeecomputersociety.org/10.1109/MCSE.2010.141> (on-line preprint)
- Tao, W.-K. and J. Simpson, 1993: The Goddard Cumulus Ensemble Model. Part I: Model description. *Terrestrial, Atmospheric and Oceanic Sciences*, **4**, 19-54.
- Tao, W.-K., C.-L. Shie, R. Johnson, S. Braun, J. Simpson, and P. E. Ciesielski, 2003: Convective Systems over South China Sea: Cloud-Resolving Model Simulations. *J. Atmos. Sci.*, **60**, 2929-2956.
- Tao, W.-K., D. Anderson, R. Atlas, J. Chern, P. Houser, A. Hou, S. Lang, W. Lau, C. Peters-Lidard, R. Kakar, S. Kumar, W. Lapenta, X. Li, T. Matsui, R. Rienecker, B.-W. Shen, J. J. Shi, J. Simpson, and X. Zeng, 2008: A Goddard Multi-Scale Modeling System with Unified Physics. WCRP/GEWEX Newsletter, Vol 18, No 1, 6-8.
- Tao, W.-K., J. Chern, R. Atlas, D. Randall, X. Lin, M. Khairoutdinov, J.-L. Li, D. E. Waliser, A. Hou, C. Peters-Lidard, W. Lau, and J. Simpson, 2009: Multi-scale modeling system: Development, applications and critical issues, *Bull. Amer. Meteor. Soc.* **90**, 515-534.
- Taft, J. R., 2001: Achieving 60 gflop/s on the production cfd code overflow-mlp. *Parallel Computing*, **27**(4):521-536.