

Large atmospheric computation on the earth simulator: The LACES project

Michel Desgagné^a, Ron McTaggart-Cowan^e, Wataru Ohfuchi^b, Gilbert Brunet^a, Peter Yau^c, John Gyakum^c, Yuki Furukawa^d and Michel Valin^a

^a*Recherche en Prévision Numérique, Meteorological Service of Canada, Quebec, Canada*

^b*Earth Simulator Centre, Yokohama, Japan*

^c*Department of Atmospheric and Oceanic Sciences, McGill University, Montreal, Quebec, Canada*

^d*Tokyo University, Japan*

^e*University of Quebec at Montreal, Canada*

Abstract. The Large Atmospheric Computation on the Earth Simulator (LACES) project is a joint initiative between Canadian and Japanese meteorological services and academic institutions that focuses on the high resolution simulation of Hurricane Earl (1998). The unique aspect of this effort is the extent of the computational domain, which covers all of North America and Europe with a grid spacing of 1 km. The Canadian Mesoscale Compressible Community (MC2) model is shown to parallelize effectively on the Japanese Earth Simulator (ES) supercomputer; however, even using the extensive computing resources of the ES Center (ESC), the full simulation for the majority of Hurricane Earl's lifecycle takes over eight days to perform and produces over 5.2 TB of raw data.

Preliminary diagnostics show that the results of the LACES simulation for the tropical stage of Hurricane Earl's lifecycle compare well with available observations for the storm. Further studies involving advanced diagnostics have commenced, taking advantage of the uniquely large spatial extent of the high resolution LACES simulation to investigate multiscale interactions in the hurricane and its environment. It is hoped that these studies will enhance our understanding of processes occurring within the hurricane and between the hurricane and its planetary-scale environment.

1. Introduction

This paper documents a meteorological study in computational fluid dynamics undertaken on a state of the art supercomputing platform. The objectives of the investigation are twofold: a) to evaluate the efficiency of the MC2 model on a large number of custom vector processors within a symmetric multiprocessor (SMP) architecture type; and b) to create a high resolution benchmark dataset on which to base diagnostic studies aimed at improving our understanding of complex hurricane lifecycles.

The field of computational fluid dynamics (CFD) has been a heavy user of computing resources since its inception in the 1960's. Over the last four decades, rough adherence to Moore's law has ensured that that the increasing technological needs of CFD have been met by faster and more efficient computing platforms.

Still, advanced geophysical models such as those used in numerical weather prediction (NWP) are capable of consuming more resources than are available in most operational centers. The need for atmospheric models to accurately resolve important atmospheric features and processes at horizontal scales on the order of a kilometer and vertical scales on the order of one hundred meters has traditionally led to a division of the problem into synoptic scale modeling systems (grid spacings between 10 km and 100 km), and large eddy simulation (LES) systems (grid spacings between 1 m and 100 m).

The challenges of physical parameterization are largely avoided by LES modeling systems in which virtually all important structures and processes are resolved. The penalty associated with such fine resolution is a rapid increase in computational load since both the horizontal grid spacing and the stepping time interval must be decreased dramatically in these sys-

tems. This operational challenge, coupled with uncertainties in the initial state of the atmosphere that render deterministic interpretation of the fine-scale structures in LES results questionable, makes forecasting on LES scales impractical at the current time on the sizes of domains required for numerical forecasting guidance.

The synoptic scale modeling systems commonly used in operational NWP to provide numerical guidance to weather forecasters use a variety of parameterization schemes to represent the effects of unresolved subgrid scale processes on the resolved flow. Each of these physically-based parameterizations addresses a single unresolved process including: moist convection (e.g. Grell [11], Kuo [17], Kain and Fritsch [14]), cloud microphysical processes (e.g. Ferrier [9], Ovtchinnikov and Kogan [23], Milbrandt and Yau [19]), radiative transfers (see review by Stephens [27]), exchanges with the surface (e.g. Blackadar [3], Noilhan [22], Verseghy [29]), and others. The accuracy of each parameterization scheme is limited by both computational constraints and our understanding of the process that it represents. For example, early cloud microphysical parameterization schemes produced rain by condensing all water vapor at humidities above an empirical threshold (“dump-bucket” schemes). Some of the most advanced schemes now create discrete size spectra for water and ice particles and treat each size bin as an independent tracer that undergoes complicated interactions with each other bin [23]. The extreme computational demands of such a scheme have led to the development of bulk schemes that approximate the size spectrum of predetermined water and ice categories using continuous functions [19]. Such balance between computational demands and physical representativeness must constantly be readjusted as resources become available. Pragmatic problems such as this are compounded by the fact that different levels of parameterization are required for different resolutions within the synoptic scale modeling system paradigm. Since a model running with a 10 km grid spacing resolves many more features than the same model with a 100 km spacing, parameterizations not applied judiciously can lead to a potential double-counting of semi-resolved processes.

Self-nesting of model runs is the standard method for increasing the resolution of a synoptic scale modeling system (see, for example, Fig. 1). Outputs from an integration with a larger grid spacing are used to provide lateral boundary conditions to a higher resolution simulation, which can then be used to nest to even higher resolution. With each decrease in grid spacing (a spacing ratio of 1:5 is an accepted standard), the

time step is decreased and the nesting domain is made smaller in response to computational load constraints. The smaller size of the inner nested, high resolution domains implies that the boundaries of the high resolution grids are closer to the features of interest than they are on the outer domains used to drive the inner grids. Contamination of the results are therefore possible through spurious interactions with the nearby lower-resolution horizontal boundaries. Similarly, the features of interest may propagate out of the high resolution domain over the integration period, making diagnostic studies difficult. The two-way nesting technique mitigates the problem of boundary contamination by allowing the results on the inner domains to influence the outer domain simulations. This reduces gradients across the boundaries, but does not eliminate this aspect of the nesting problem entirely. A moving-grid solution is often implemented to address the high resolution domain size problem: the inner grids track a feature of interest and move through the outer grids over the course of the simulation [16]. Although this method allows for the high resolution simulation of a feature over an extended period of time, it introduces additional computational problems and does not allow for a broad scale view of the event from a high resolution perspective.

The simulation of a meteorologically interesting event at high resolution over a domain large enough to capture both the fine scale details of the system and the broad structures of the surrounding flow is the purpose of the current study. The resulting dataset will be used for diagnostics of the event itself [Hurricane Earl (1998)] and for further investigation into optimal parameterization and nesting techniques. The generation of a high resolution benchmark run was impossible before the creation of the Japanese Earth Simulator Center (ESC) in 2002. The ESC controls the Earth Simulator (ES) highly parallel vector supercomputer, comprising 640 NEC SX-6+ nodes – each of which contain 8 processing elements (PE). The peak theoretical performance for the ES is 40 Tflop with a total main memory capacity of 10 TB. For detailed information on the design of the ES, the reader is referred to Sato [25].

A long-term collaborative effort between scientists from the ESC, McGill University and Recherche en Prévision Numérique (RPN) has now performed high resolution, large domain simulation of the full life cycle of Hurricane Earl (August–September 1998) with the Canadian Mesoscale Compressible Community (MC2) Model [10]. The storm presents a challenge to high resolution numerical weather prediction systems as it tracks more than 6000 km from the Gulf of Mexico

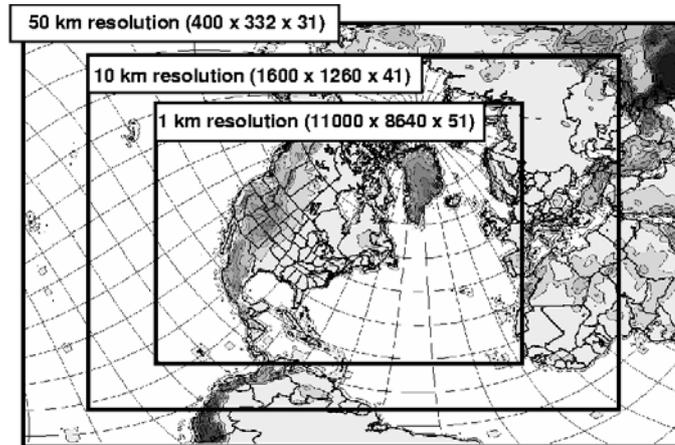


Fig. 1. Computational domains for the LACES project.

across the North Atlantic Ocean during its 8 day life-cycle.

Hurricane Earl, a Category 1 storm [26] at landfall in the Florida Panhandle, reached its peak intensity near the Canadian Maritimes during its extratropical transition (ET) into a midlatitude system (McTaggart-Cowan et al. [18] present a review of the later stages of Hurricane Earl's lifecycle). Large storm structure and intensity variations associated with ET events such as Hurricane Earl present a unique challenge to numerical models. Errors in the strength or nature of the flow in and around the system propagate rapidly upscale and degrade numerical forecast guidance over a broad area both upstream and downstream of the system [13]. As noted above, the primary meteorological objective of this collaborative project (LACES – Large Atmospheric Calculation on the Earth Simulator) is to create a high resolution benchmark dataset on which to base diagnostic studies aimed at improving our understanding of the lifecycle of a complex hurricane.

With a horizontal grid spacing of 1 km on a domain extending from western Pacific to eastern Europe, LACES requires the unique and extensive computing resources currently available only at a few centers like the Earth Simulator Center. The MC2 model was installed on the Earth Simulator in November 2003. After demonstrating the basic performances and scalability of the model on a large number of ES' processors, production started in September 2004. The life cycle of Hurricane Earl between 1800 UTC 31 August and 1200 UTC 7 September has now been simulated over a large domain. A database containing the results of the simulation is being built at RPN, and diagnostic studies have started.

This paper begins with a description of the computational problem undertaken as the LACES project in Section 2. Section 3 presents an analysis of the MC2 model's structure and performance on the ES system. Details of the experimental design and production phase are covered in Section 4. Preliminary validation of the results of the simulation are presented in Section 5. The study concludes with summary and discussion in Section 6.

2. Computational goals of the LACES project

The preliminary goal of LACES is to produce a credible high resolution simulation of Hurricane Earl's full lifecycle over a very large domain (Fig. 1). In order to capture both the tropical and extratropical-redevelopment phases of the storm, 7.25 days of simulation with a 1 km grid spacing on a $11000 \times 8640 \times 51$ domain are required. This reference simulation will be used to compare with various lower resolution simulations and to improve our understanding of hurricane dynamics, ET, multiscale hurricane/background flow interactions, and predictability issues associated with tropical cyclones throughout their lifecycles. Of course this represents a major computational effort that could only be performed on the Earth Simulator at the time of the experiment. Over the course of the simulation, the MC2 used over 75% of the total resources of the system for eight full days of computation.

This project also represents a challenge for the MC2 modeling system itself considering that the computational load is distributed on 3960 vector processors (495 out of the 640 ES nodes) [7]. Inter-processor commu-

nications become a potential issue with such a large number of processors [28]. LACES also represents a rigorous test for the numerics of the elliptic solver considering that order 10^9 equations are solved at once by an iterative three dimensional pressure solver based on FGMRES [24,28]. Input and output (I/O) are also very important issues that must be treated carefully. Even the simple display of full meteorological fields on this size of horizontal plane is far beyond the limit of current display technology.

3. MC2 code structure, scalability and performance

The MC2 model is a production quality weather forecast model widely used at Environment Canada and in Canadian Universities. It solves a full set of Euler equations using a three-time level semi-implicit semi-Lagrangian time integration scheme [10]. In space the model is discretized using centered finite differentiation on a $N_x \times N_y \times N_z$ staggered Arakawa ‘C’ grid [2] with constant horizontal grid spacing ($\Delta X, \Delta Y$). Equations are solved on a limited area domain for which lateral boundary conditions must be prescribed. This is usually done using data from a previous run at coarser resolution on a larger domain or using a set of analysis data. The elliptic problem that arises from the semi-implicit time discretization scheme in MC2 leads to a system of equations with a large sparse non-symmetric coefficient matrix. The problem is solved using a flexible variant of the GMRES algorithm (FGMRES) proposed by Saad [24] by which the preconditioner can be adjusted or entirely substituted during the iterative process in order to accelerate convergence. The only preconditioner used for the LACES project is a local vertical line relaxation preconditioner. Finally, a set of physical parameterization routines is called at the end of each timestep in order to model the most important processes which remain unresolved due to the choice of horizontal resolution. Time looping can be summarized as followed:

```
DO timestep=1, N
  Call RSH      ! compute right hand side of
                differential equations
  Call SL       ! perform semi-Lagrangian
                advection of RHS
  Call SOLVER   ! solve the elliptic problem
                for pressure
  Call BAC      ! backward substitution to
                obtain solutions for all other
                variables
```

```
Call          ! parameterize unresolved
PHYSICS       features
Call HORD     ! Apply numerical horizontal
                diffusion
END DO
```

The entire MC2 library is written in standard Fortran with main memory allocated dynamically to suit the many configurations the model actually allows. In terms of computing resources the most demanding algorithms are most certainly the semi-Lagrangian interpolations and the GMRES solver. An efficient single program multiple data implementation of both these algorithms is crucial to achieve good performance. Automatic vectorization is used on vector processor architecture and the global vectorization ratio will exceed 98% on most platforms. Parallel programming paradigms implemented in MC2 are twofold. First a horizontal domain decomposition technique has been implemented [28] whereby the $(N_x \times N_y)$ horizontal domain is partitioned across a $P_x \times P_y$ logical processor mesh so that a subdomain contains $(N_x/P_x) \times (N_y/P_y) \times N_z$ points. Local subdomain are surrounded by a halo region of 3 points to contain grid point data communicated by adjacent processor. Inter processor boundary exchanges are implemented using MPI [8]. Communication load is determined by the actual data flow itself and the local requirement of individual algorithm. On subdomain size of order 250×50 (for vector processors) the time spent communicating is overall 0.1% of the total execution time. It is of course possible to drastically increase this load by over decomposing a global problem size into local subdomains with too few computational grid points. Over decomposing would also cause a noticeable increase in total memory usage but this is normally not a major concern since the halo region only contains 3 grid points. Second, a ‘do loop’ parallelization technique using OpenMP [6] is implemented as an independent feature allowing a hybrid combination of both methods whenever possible. The scalability of the current OpenMP implementation is actually not sufficient mainly because single and/or critical segments are still too numerous. We therefore elected to only use the MPI implementation on the ES. No code modification was necessary to run MC2 on the ES. Only a layer of surrounding job language scripts needed to be constructed in order to control the actual batch jobs on the ES.

Before starting the production run with the MC2 model, a few tests were done in order to ensure efficiency and scalability of the code. Given the fact that the code was already vectorized to more than 98% and

was running at about 25% of the peak performance on a single PE, we elected to only test the scalability of the MPI code on a large number of processors [7]. This was done using first a scalable problem size. In order to take advantage of the full vector length and considering that most of the MC2 code is vectorized along west-east horizontal axis NI, it is best to design subdomains that have local NI as close as possible to a multiple of 256 which is the current vector length on the ES' processors. This will of course limit the size of the other horizontal dimension NJ (south-north axis) if one is to make the subdomains as small as possible. We found that a subdomain shape of 500×50 achieves a good compromise between computational load, communication and memory size to fit on a single processor. This shape offers almost two full vector lengths and yet is not really demanding in terms of inter-processor communications despite the narrow rectangular shape.

One can now establish a time to solution running the model on a single subdomain for a certain number of timesteps. We then construct various global domain sizes by assembling together those unit 500×50 subdomains. Running the model for the same number of timesteps using the same number of processors as there are subdomains should in principle yield the same time to solution. This of course is not entirely true for a global domain composed of 1 to about 9 subdomains as communication patterns set in and consequently alter the scalability slightly. From then on however we were able to demonstrate almost perfect scalability for up to 1120 processors. This exercise also helped to put aside one major concern we had about the convergence of the iterative solver as the global problem size becomes larger and larger. It turns out that the number of outer iterations performed by GMRES barely increases by 1 to 3 as the global problem size increases from 500×50 to 5000×5600 . From this point we were already confident that the code would scale to 3960 processors on the targeted global domain of 11000×8640 .

Scalability tests were also performed on fixed size problems in order to satisfy requirements of the ESC regarding the attribution of resources. To that effect the maximum number of processors 'M' the ESC will allow an application to use is based on the parallelization ratio α and is given by:

$$M \leq \frac{2 - \alpha}{1 - \alpha} \quad (1)$$

with

$$\alpha = \frac{m-1}{m} \frac{T_n - T_m}{T_n - \frac{n-1}{n} T_m} \quad (2)$$

where:

T_n = execution time on n processors

T_m = execution time on m processors

($m > n$)

We achieved early on enough scalability to be allowed 140 nodes. With these resources we were able to fit a global problem size of $5000 \times 5040 \times 51$ (horizontal x-dimension \times horizontal y-dimension \times vertical dimension) with subdomains of size $500 \times 45 \times 51$ when running on the full 140 nodes using a 10×112 processor topology. This produced a wall clock time to solution of 2366 sec. In order to minimally affect the vector length, the same run was then performed using a 10×56 processor topology. This yielded a wall clock time to solution of 4553 sec. Using these results to compute the parallelization ratio α will yield a value of 99.99% and leads to $M < 13690$. This is well above the ES total computing resources and it simply means that the code scales well enough for our purpose on the ES. The project was hence allowed the required 495 needed for production. Preliminary tests on the full scale 1 km domain with dimension $11000 \times 8640 \times 51$ using a 22×180 processor topology were performed in May 2004. The MC2 model version 4.9.7 globally sustained 2.5 GFlops per processor for an aggregate speed of 9.9 TFlops. The dynamic kernel alone sustained 3.1 GFlops per processor for an aggregate performance of 12.3 TFlops. The total memory required to fit such a large problem is 1.9 GB per processor for a total of 7500 GB.

4. Experimental design and production

As noted in Section 2, the preliminary goal of the LACES project is to produce a simulation over a very large domain with a 1 km grid spacing which covers the tropical and extratropical-redevelopment phases of Hurricane Earl's lifecycle. The simulation strategy is a triple self-nested grid system that starts with a 50 km (grid spacing) outer domain driven by Canadian Meteorological Centre (CMC) analysis which is used to drive a 10 km domain which in turns is used to drive the targeted 1 km domain (Fig. 1). These analysis fields are chosen for the LACES study because the MC2 model has been shown to perform well using such a nesting strategy for simulations of Hurricane Earl [18]. Although a moving nest for the innermost domains would allow for computations over a smaller area, the dynamical consistency of such an approach is problematic given the need for interpolation as the grid moves.

Conservative moving-nest techniques have been developed for meteorological applications [16]; however, the multi-scale and wave diagnostics planned for the analysis stage of the LACES project require a uniform high resolution grid over a broad area. The timesteps are 240, 60 and 6 seconds, respectively for the 50, 10 and 1 km domains, designed to keep the Courant-Friedrich-Lewy (CFL) criterion approximately unity. The CFL criterion is not a strict stability limitation on the semi-Lagrangian advection of the MC2 model, but ensures accuracy by preventing air parcels from advecting further than the length of the grid box in a single time step. A Kain-Fritsch cumulus parameterization scheme [14] is used for the 50 and 10 km domain. The precipitation is entirely explicit at 1 km using an advanced micro-physics package [15]. All runs are performed with an explicit horizontal hyper-diffusion scheme of the form

$$\frac{\partial F}{\partial t} = K \nabla^n F \quad (3)$$

with $n = 6$ and the diffusion coefficient K chosen in such a way that the scheme will remove 10% of the $2\Delta x$ signal every timestep (i.e. 10% suppression of features having wavelengths of twice the horizontal grid spacing). With a 6 seconds timestep we get a ratio a bit larger than unity between wall clock computing time and simulation time (i.e. it takes a bit more than one hour of wall clock computing time to simulate one hour). We therefore expected to use over eight days of computation on 495 nodes to complete the task.

Because of predictability concerns over such a long integration period, particularly for the ET and re-intensification of Hurricane Earl, the model is reconnected to the analysis at the end of the tropical phase and at the end of the transition phase (Fig. 2). A 6 hour time overlap is required to allow the dynamical and physical processes of the model to develop the flow in the interior of the domain that is consistent with the enhanced resolution of the inner grid in comparison to the outer driving model. This ‘‘spin up’’ time is typical for nested limited area models (LAMs).

The production runs started in September 2004. The computation was spread over 75% of the total computing resources of the ESC and the execution time per job was limited to 4 hours. We therefore proceeded with a restart technique whereby at the end of each computation period every processor element (PE) writes up to disk the current state of a large portion of its core memory. As for the model outputs themselves, this I/O is done using file systems locally attached to every node and therefore limits the bottleneck effect. Any single

job therefore starts with a stage-in phase to bring the restart files and the lateral boundary conditions from the front-end machine to the proper ES nodes. The model is then launched. Every PE will read its own restart file and then compute for 4 wall clock hours before writing a restart file again. The job ends with a stage-out phase to bring the restart files back to the front-end along with the model output files. Given the fact that the restart files themselves total close to 2 TB of data, a single job typically takes about 12 hours to complete. The waiting period on the ES’ batch system queue is highly variable. The extratropical phase was completed in June 2005.

5. Validation and diagnostics

The primary verification parameters for hurricane simulation are storm intensity and track. Although track forecasting skill in numerical models has increased dramatically over the last decade, quantitative estimates of hurricane genesis and intensity remain problematic. Although higher resolution operational forecast models have been shown to improve intensity guidance [1], limited domain size and possible boundary problems still reduce the effectiveness of these integrations. As noted in the introduction, the primary meteorological goal of the LACES project is to create a high resolution benchmark dataset on which to base future studies. It is therefore important for the integration to explicitly resolve important forcings and features (for example, vortex Rossby waves [20], vortical hot towers [12] and other near-core structures) over a domain that is sufficiently large to minimize the deleterious impact of lateral boundary conditions throughout the simulation. The extremely large inner nested grid implies that the size of the model’s state space is maximized (given computational constraints), as are the number of degrees of freedom throughout the simulation. The evolution of the broad area of the inner domain as a region that is largely free of artificial external forcings will have two important implications: 1) it will increase the risk that the model’s trajectory through state space will diverge from that of the observed system; and 2) it will allow the model to explicitly capture small scale forcings involved during the storm’s lifecycle, and to accurately simulate their interactions both with each other and with their environment. Errors in the LACES simulation that arise from item (1) will be of interest because the planned testing of improved physical packages on subdomains of the inner

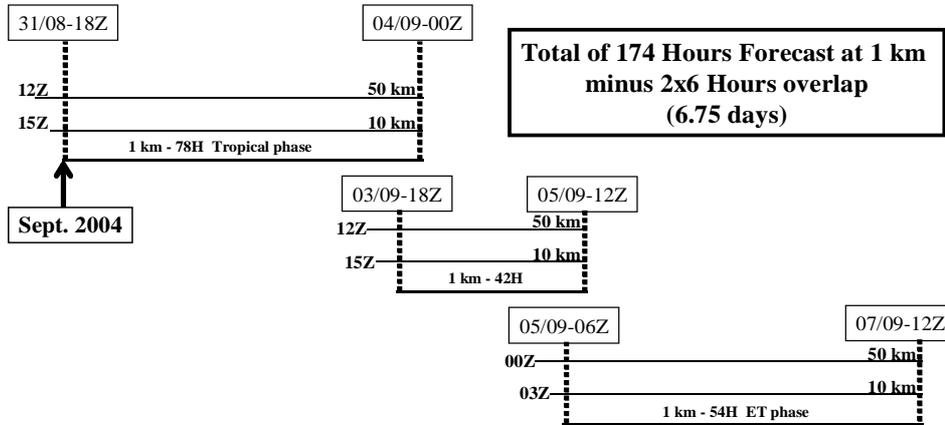


Fig. 2. Time strategy for the simulations of the full lifecycle of Hurricane Earl on the ES.

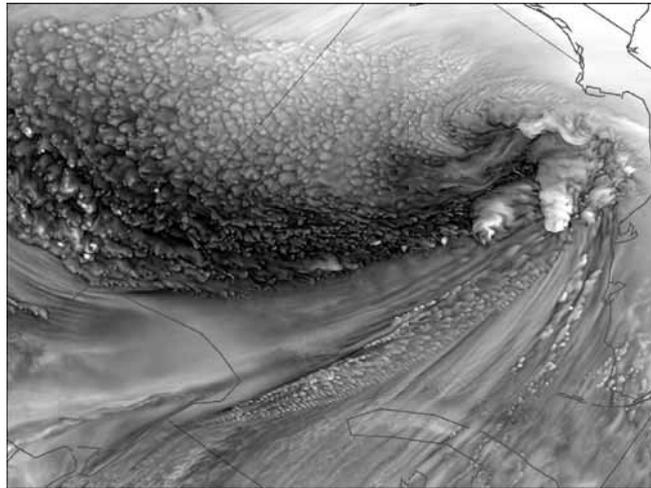


Fig. 3. Simulated 45 h near surface (2 m) specific humidity valid 1500 UTC 2 September 1998. Hurricane Earl is at this stage fully developed in the LACES simulation and is about to landfall in western Florida State.

LACES grid will allow for a minimization of this error mode in future simulations and operational forecasts. Diagnostic studies focusing on item (2) will allow for the evaluation of scale interactions and high resolution processes in a dynamically consistent environment that is as free from boundary contamination as currently possible in LAMs.

Standard hurricane intensity verification consists of a comparison between the modeled minimum sea level pressure (SLP) and the SLP archived in the Best Track dataset produced by the National Hurricane Center (NHC) for each named storm. As shown in Fig. 4, the LACES simulation captures the early intensification of Hurricane Earl over the first 48 h of integration (between 1800 UTC 31 August and 1800 UTC 2 August) from a tropical depression to a Category 1 Hurri-

cane [26]. The track of the modeled system has both a right-of-track and a fast bias throughout the simulation (Fig. 5) that appear to result directly from a discrepancy between the Best Track position of the storm at 1800 UTC 31 August (22.4°N, 93.8°W) and the analyzed position used for model initialization (25.0°N, 92.1°W). Especially during the weak initial stages of tropical cyclogenesis, the center of the developing system can be difficult to locate precisely, leading potentially to track biases similar to those shown in Fig. 5. The along-track position error also leads to inconsistencies between the SLP traces of the simulated storm and the Best Track record (Fig. 4). The simulated system interacts too early with the Florida Peninsula, resulting in SLP increases and a pressure error of up to 10 hPa by 0000 UTC 3 September (Best Track landfall). Because

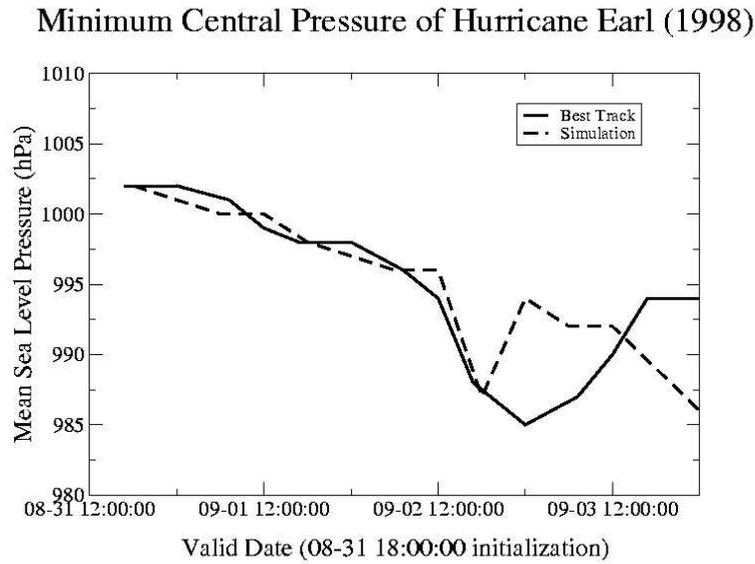


Fig. 4. Time series of Best Track (solid line) and simulated (dashed line) minimum SLP at the storm center.

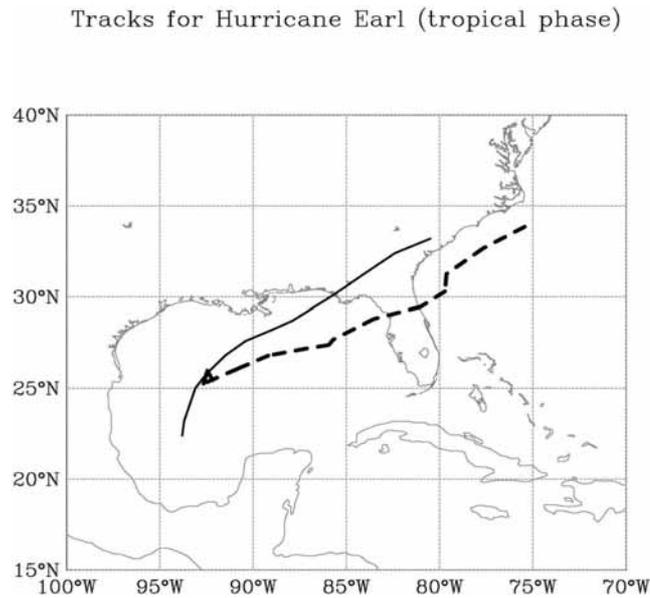


Fig. 5. Track of Hurricane Earl from the Best Track archive (black solid line) and simulated (red dashed line). Both tracks start at 1800 UTC 31 August and end 0000 UTC 4 September.

quantitative track forecast skill does not rely heavily on model resolution (model track verification statistics compiled over the last decade by the NHC show that global models running at low resolution may outperform higher resolution regional models on any given case), the failure of the high resolution LACES simulation to correct the initialization-induced track bias is

not unexpected. However, the shift in the vortex location does not appear to adversely affect the intensity (Fig. 4) and structural evolution of the system. The latter is shown in detail in Fig. 3 as represented by the 45 h near-surface (325 m above mean sea level) specific humidity over a limited subdomain of the model grid. What makes this simulation unique is the generation of

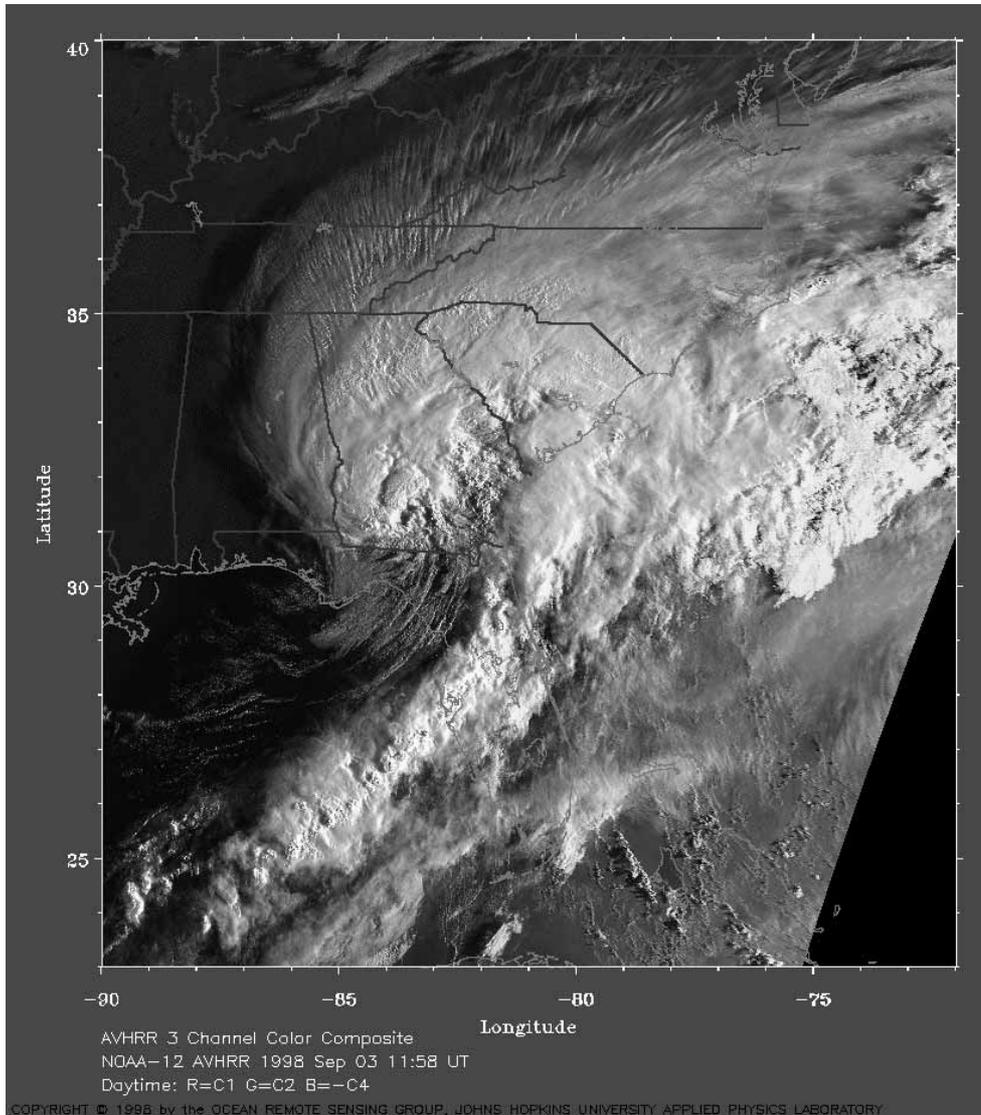


Fig. 6. Advanced Very High Resolution Radiometer (AVHRR) National Oceanic and Atmospheric Administration (NOAA-12) satellite image of Hurricane Earl valid at approximately 1200 UTC 3 September. Lighter colors generally indicate higher cloud tops (colder temperatures), but the low angle sunlight from the east casts shadows to the west of convective tower.

such high resolution information over an area that covers much of the western hemisphere (Fig. 1). Previous studies have only been able to resolve such fine scale features over a very limited area and therefore disallow explicit interaction with the large scale flow.

Although Hurricane Earl reaches minimal Category 2 intensity (near surface winds exceeding 85 kt) at 1800 UTC 2 September, 48 h after the initialization of the LACES simulation, it never develops a classic symmetric hurricane structure. As shown in Fig. 6 by 1200 UTC 3 September, just 12 h after landfall, Earl has a distinctly asymmetric structure as it crosses the south-

eastern United States. The rapid structural evolution of the system over the course of the 78 h LACES simulation – from weak tropical depression to Category 2 hurricane to asymmetric cyclone – poses a unique challenge to the MC2 modeling system. Figure 7 shows the 66 h cloud top temperatures from the LACES simulation, verifying at the same time as the satellite image in Fig. 6 (approximately 1200 UTC 3 September 1998). The domain of Fig. 7 has the same dimensions as Fig. 6, but is shifted to the east to account for the bias in the simulation track noted earlier in this section. The structure of the storm, as diagnosed by the cloud patterns in

Cloud Top Temperature

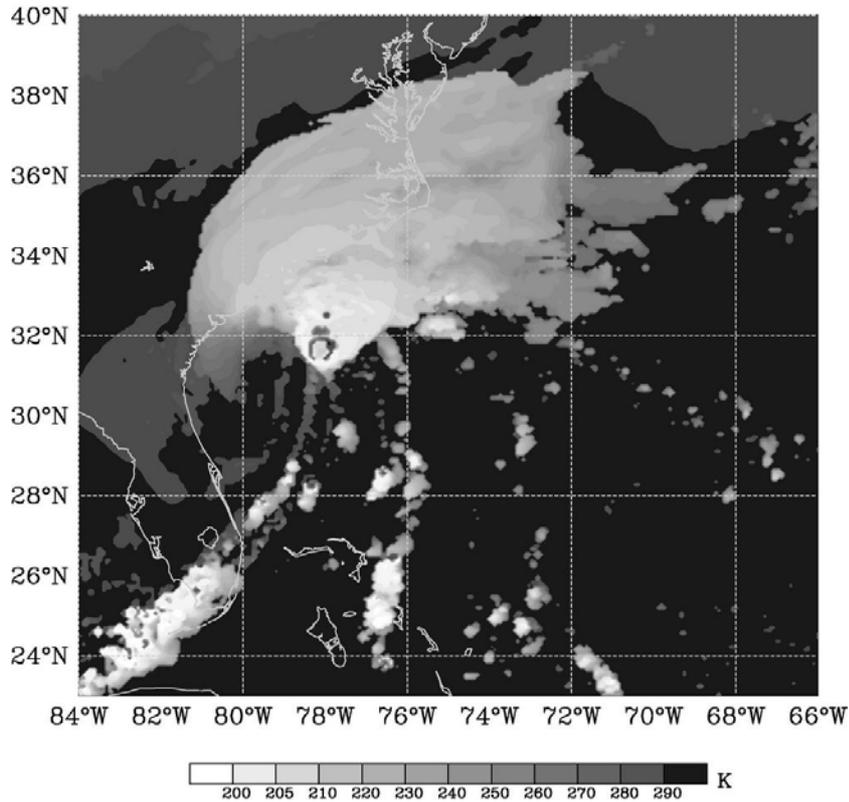


Fig. 7. Cloud top temperatures from the LACES simulation (66 h) valid at 1200 UTC 3 September for comparison with Fig. 6. Temperatures in Kelvin as indicated on the color bar beneath the figure. Note that the domain is the same size as that of Fig. 6, but is shifted 6E to account for the right-of-track bias in the simulation.

Fig. 7, closely resembles that of the observed system (Fig. 6). The MC2 model successfully simulates the “delta” rain pattern to the north of the center of circulation (near 31.5°N, 78°W in Fig. 7), and generates an intense band of convection that extends across southern Florida.

Despite the right-of-track bias of the storm in the LACES simulation, the structure and evolution of Hurricane Earl throughout the tropical phase of its lifecycle are well simulated by the model. A challenging measure of the model’s performance lies in the analysis of the quantitative precipitation forecast produced over the course of the simulation. The storm-total precipitation [computed from model fields verifying between 1200 UTC 2 September (42 h) to 0000 UTC 4 September (78 h, the end of the simulation)] is shown in Fig. 8 for comparison with the observed storm-total precipitation

from the Unified Precipitation Dataset¹ (UPD, computed from daily totals for the 1200 UTC 2 September to 1200 UTC 4 September period) shown in Fig. 9. The discrepancy in the temporal coverage between Figs 8 and 9 is a result of the 24 h accumulation period used to generate the UPD dataset and is not expected to result in significant errors in this subjective comparison of the accumulated precipitation fields. The spatial coverage of the rain gauge and radar-derived UPD is limited to the continental United States, making comparison of Figs 8 and 9 over ocean waters impossible. Given the large values of rainwater accumulation (beyond 100 mm in both the model and observed fields

¹Climate Prediction Center United States Unified Precipitation data provided by the NOAA-CIRES Climate Diagnostics Center, Boulder, Colorado, USA, from their Web site at <http://www.cdc.noaa.gov/>.

Model Precipitation

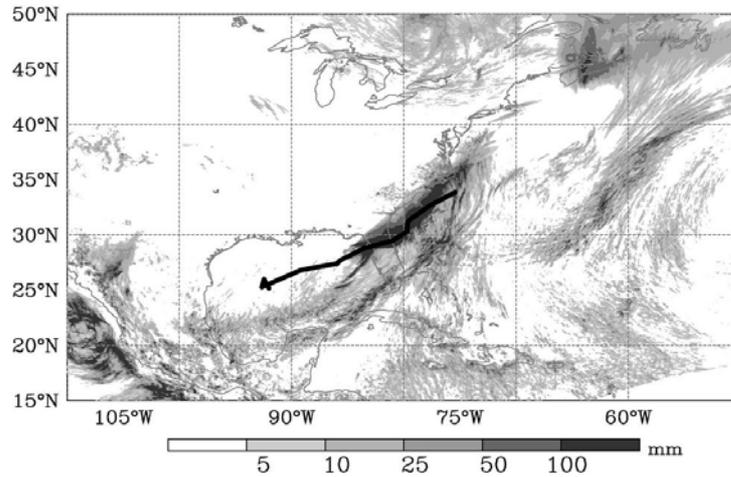


Fig. 8. Accumulated precipitation (in mm, as indicated on the color bar) from the LACES simulation between 1200 UTC 2 September (42 h) and 0000 UTC 4 September (78 h). The track of Hurricane Earl in the simulation is shown with the solid black line.

UPD Analyzed Precipitation

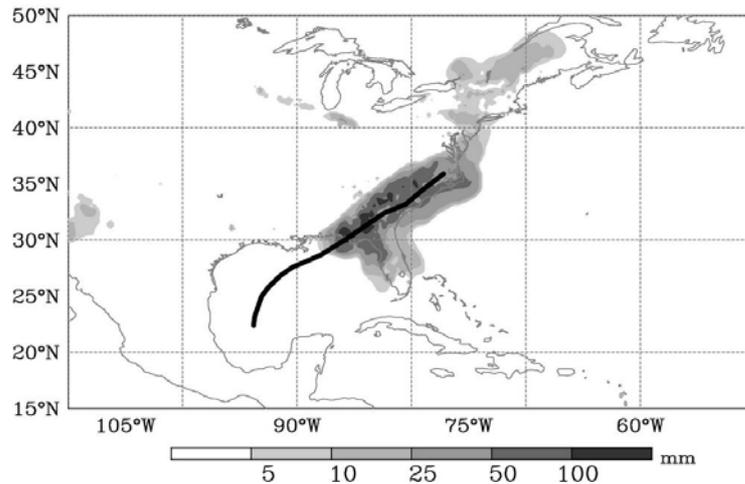


Fig. 9. Observed accumulated precipitation (in mm, as indicated on the color bar) from the UPD (see footnote in text for source). The Best Track of Hurricane Earl is shown with the solid black line. In comparison with Fig. 8, note that the UPD covers only land, and is of highest quality over the continental United States where rain gage data and full radar coverage make analyses more credible.

over the two day period), the primary forcing for this event is undoubtedly the hurricane itself, acting in relative isolation from any other atmospheric features of significance. The error in the simulated track of Earl is once again evident, but the spatial extent and magnitude of the precipitation field generated by the model compare well with the observed values. The location of the precipitation maximum with respect to the storm track (shown by the thick black line in Figs 8 and 9) is

also consistent between model and observations, suggesting that the structures and forcings in the hurricane are well modeled in a storm-relative sense.

6. Summary and conclusions

The LACES project represents a major collaborative initiative between Canadian and Japanese meteorolog-

ical services and researchers, focusing on the high resolution simulation of Hurricane Earl (1998). The preliminary computational goal of the project, achieved in August 2005, was the fully-scalable implementation of the MC2 modeling system on the ES computers. Despite the large domain size of $11000 \times 8640 \times 51$, the MC2 model achieved a parallelization ratio of 0.9999:1 and occupied over 75% of the ESC resources to run at a rate of 9.9 TFlop. The 5.2 TB of unprocessed output from the LACES simulation has been moved from the ESC to RPN, where it is housed in a database that has been made available to participating researchers.

Work towards the specific scientific goals of the LACES project has only begun since the completion of the production stage of the simulation. Preliminary diagnostics have been presented here, and indicate that the LACES simulation of the tropical phase of Hurricane Earl's lifecycle successfully reproduces many of the observed storm structures. A right-of-track bias is noted in the simulation results, but the generally isotropic nature of the tropical environment appears to minimize errors arising from the mis-location of the storm at later times in the simulation. The reinitialization of the model starting at 1800 UTC 3 September using nested fields derived from the 1200 UTC 3 September analysis will ensure that subsequent (transition and reintensification) stages of Earl's lifecycle will not be adversely affected by the track error that develops over the course of this tropical-stage simulation. Verification of these stages of the LACES simulation is under way and will form the bulk of a future study to be presented in the meteorological literature.

Advanced post-processing of the data, including energy budget, empirical normal mode [4,5] and dynamic tropopause analyses [21] have begun at RPN and academic institutions involved with the LACES project. Modeling studies aimed at improving physical parameterizations over a wide range of horizontal scales using the LACES dataset as a benchmark are planned. The large extent of the high resolution data, the unique quality of this dataset, will be fully exploited in these studies. Multiscale interactions between the vortex and its immediate environment, and the local environment and the hemispheric-scale flow will be diagnosed at high resolution for – to our knowledge – the first time. It is hoped that these investigations of the data generated by the LACES project will yield new insights into the role of fine-scale structures in the development and evolution of tropical systems. In turn, an improved understanding of these usually unresolved processes may assist in the operational prediction of hurricanes in real time.

References

- [1] M.A. Bender and I. Ginis, Real-Case Simulations of Hurricane – Ocean Interaction Using A High-Resolution Coupled Model: Effects on Hurricane Intensity, *Mon. Wea. Rev.* **128** (2000), 917–946.
- [2] R. Benoit, M. Desgagne, P. Pellerin, S. Pellerin, Y. Chartier and S. Desjardins, The Canadian MC2: A semi-Lagrangian, semi-implicit wide-band atmospheric model suited for fine-scale process studies and simulation, *Mon. Wea. Rev.* **125** (1997), 2382–2415.
- [3] A.K. Blackadar, Modeling the nocturnal boundary layer. *Proc. Third Symp. on Atmospheric Turbulence, Diffusion and Air Quality*, Boston, Amer. Meteor. Soc., 1976, 46–49.
- [4] G. Brunet, Empirical normal-mode analysis of atmospheric data, *J. Atmos. Sci.* **51** (1994), 932–952.
- [5] Y. Chen, G. Brunet and M.K. Yau, Asymmetric structures in a simulated landfalling hurricane, *J. Atmos. Sci.* **60** (2003), 2294–2312.
- [6] D. Clark, Focus: OpenMP: a parallel standard for the masses, *IEEE Concurrency* **6** (1998), 10–12.
- [7] M. Desgagné, S.J. Thomas and M. Valin, Performance of MC2 and the ECMWF IFS forecast model on the Fujitsu VPP700 and the NEC SCX4M, *Scientific Programming* **8** (2000), 23–30.
- [8] J.J. Dongarra and D.W. Walker, MPI: A standard message passing interface, *Supercomputer* **12** (1996), 56–68.
- [9] B.S. Ferrier, A double-moment multiple-phase four-class bulk ice scheme. Part I: Description, *J. Atmos. Sci.* **51** (1994), 249–280.
- [10] C. Girard, R. Benoit and M. Desgagné, Finescale topography and the MC2 dynamics, *Mon. Wea. Rev.* **133** (2005), 1463–1477.
- [11] G. Grell, Prognostic evaluation of assumptions used by cumulus parameterizations, *Mon. Wea. Rev.* **121** (1993), 764–787.
- [12] E.A. Hendricks, M.T. Montgomery and C.A. Davis, On the role of vortical hot towers in hurricane formation, *J. Atmos. Sci.* **61** (2003), 1209–1232.
- [13] S. Jones and coauthors, The extratropical transition of tropical cyclones: forecast challenges, current understanding, and future directions, *Weather and Forecasting* **18** (2003), 1052–1092.
- [14] J.S. Kain and J.M. Fritsch, A one-dimensional entraining/detraining plume model and its application in convective parameterization, *J. Atmos. Sci.* **47** (1990), 2784–2802.
- [15] F.Y. Kong and M.K. Yau, An explicit approach to microphysics in MC2, *Atmos.-Ocean* **35** (1997), 257–291.
- [16] Y. Kurihara, G.J. Tripoli and M.A. Bender, Design of a movable nested-mesh primitive equation model, *Mon. Wea. Rev.* **107** (1979), 239–249.
- [17] H.L. Kuo, Further studies of the parameterization of the influence of cumulus convection on large scale flow, *J. Atmos. Sci.* **31** (1990), 1232–1240.
- [18] R. McTaggart-Cowan, J.R. Gyakum and M.K. Yau, Sensitivity testing of extratropical transitions using potential vorticity inversions to modify initial conditions: Hurricane Earl case study, *Mon. Wea. Rev.* **129** (2001), 1617–1636.
- [19] J. Milbrandt and M.K. Yau, A multimoment bulk microphysics parameterization. Part II: a proposed triple moment closure and scheme description, *J. Atmos. Sci.* **62** (2005), 3065–3081.
- [20] M.T. Montgomery and R.J. Kallenbach, A theory for vortex Rossby waves and its application to spiral bands and intensity changes in hurricanes, *Quart. J. Roy. Meteor. Soc.* **123** (1997), 435–465.

- [21] M.C. Morgan and J.W. Neilson-Gammon, Using tropopause maps to diagnose midlatitude weather systems, *Mon. Wea. Rev.* **126** (1998), 2555–2579.
- [22] J. Noilhan and S. Planton, A simple parameterization of land surface processes for meteorological models, *Mon. Wea. Rev.* **117** (1989), 536–549.
- [23] M. Ovtchinnikov and Y.L. Kogan, An investigation of ice production mechanisms in small cumuloform clouds using a 3D model with explicit microphysics. Part I: model description, *J. Atmos. Sci.* **57** (2000), 2989–3003.
- [24] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, *SIAM Journal on Scientific Computing* **14** (1993), 461–469.
- [25] T. Sato, The current status of the Earth Simulator. *Journal of Earth Simulator* **1** (2004), 6–7, available at <http://www.es.jamstec.go.jp/esc/images/journal200404/index.html>.
- [26] R.H. Simpson, The hurricane disaster potential scale, *Weatherwise* **27** (1974), 169–186.
- [27] G.L. Stephens, The parametrization of radiation for numerical weather prediction and climate models, *Mon. Wea. Rev.* **112** (1984), 826–867.
- [28] S.J. Thomas, A.V. Malevsky, M. Desgagné, R. Benoit, P. Pel-lerin and M. Valin, Massively parallel implementation of the Mesoscale Compressible Community model, *Parallel Computing* **23** (1997), 2143–2160.
- [29] D. Verseghy, CLASS – A Canadian land surface scheme for GCMs. I: Soil model, *Int. J. Climatol.* **11** (1991), 111–133.



Hindawi

Submit your manuscripts at
<http://www.hindawi.com>

