

A scalable high-order discontinuous Galerkin method for global atmospheric modeling*

Hae-Won Choi^{a†}, Ramachandran D. Nair^{a‡} and Henry M. Tufo^{a§}

^aScientific Computing Division, National Center for Atmospheric Research (NCAR),
1850 Table Mesa Drive, Boulder, CO 80305, USA

A conservative 3-D discontinuous Galerkin (DG) baroclinic model has been developed in the NCAR High-Order Method Modeling Environment (HOMME) to investigate global atmospheric flows. The computational domain is a cubed-sphere free from coordinate singularities. The DG discretization uses a high-order nodal basis set of orthogonal Lagrange-Legendre polynomials and fluxes of inter-element boundaries are approximated with Lax-Friedrichs numerical flux. The vertical discretization follows the 1-D vertical Lagrangian coordinates approach combined with the cell-integrated semi-Lagrangian method to preserve conservative remapping. Time integration follows the third-order SSP-RK scheme. To valid proposed 3-D DG model, the baroclinic instability test suite proposed by Jablonowski and Williamson is investigated. Parallel performance is evaluated on IBM Blue Gene/L and IBM POWER5 p575 supercomputers.

1. INTRODUCTION

The future evolution of the Community Climate System Model (CCSM) into an Earth system model will require a highly scalable and accurate flux-form formulation of atmospheric dynamics. Flux form is required in order to conserve tracer species in the atmosphere and accurate numerical schemes are essential to ensure high-fidelity simulations capable of capturing the convective dynamics in the atmosphere and their contribution to the global hydrological cycle. Scalable performance is necessary to exploit the massively-parallel petascale systems that will dominate high-performance computing (HPC) for the foreseeable future.

The High-Order Method Modeling Environment (HOMME) [5], developed by the Scientific Computing Section at the National Center for Atmospheric Research (NCAR), is a vehicle to investigate using high-order-element-based methods to build conservative and accurate dynamical cores. HOMME employs the spectral element (SE) methods on a cubed-sphere tiled with quadrilateral elements, can be configured to solve the shallow water or the dry/moist primitive equations, and has been shown to efficiently scale to 32,768 processors of an IBM Blue Gene/L (BG/L) system [16]. Nevertheless, a major disadvantage of the SE atmospheric model is that it is not inherently conservative. For climate and atmospheric applications, conservation of integral invariants such as mass and total energy is of significant importance. To resolve these issues, we recently included the DG atmospheric models to support HOMME framework.

In this paper we discuss our extension of the HOMME framework to include a 3-D DG option as a first step towards providing the atmospheric science community a new generation of atmospheric general circulation models (AGCMs). The DG method [2], which is a hybrid technique combining the finite element and finite volume methods, is inherently conservative and shares the same computational advantages as the SE method such as scalability, high-order accuracy, spectral convergence, and thus is an ideal candidate for climate modeling. The DG method is employed on a quadrilateral mesh of elements using a high-order nodal basis set of orthogonal Lagrange-Legendre polynomials with Gauss-Lobatto-Legendre (GLL) quadrature points. Time integration follows the strong stability-preserving Runge-Kutta (SSP-RK) scheme of Gottlieb *et al.* [4]. The globe is based on the singularity-free cubed-sphere geometry

*This paper is supported by the **DOE-SciDAC** program under award **#DE-FG02-04ER63870**.

†Corresponding Author. Email: haewon@ucar.edu.

‡Email: rnair@ucar.edu.

§Email: tufo@ucar.edu.

introduced by [13]. Parallelism is effected through a hybrid MPI/OpenMP design and domain decomposition through the space-filling curve approach described in [3]. Our work extends earlier efforts [3,10–12] in several important ways: first, we develop a scalable conservative 3-D DG-based dynamical core based on the hydrostatic primitive equations; second, we employ the vertical Lagrangian coordinate approach, developed by Starr [15] and later generalized by Lin [8]; and finally, we apply the 1-D cell-integrated semi-Lagrangian method [9] to preserve conservative remapping.

2. CONSERVATIVE DISCONTINUOUS GALERKIN MODEL

Inherently conservative numerical schemes are of fundamental importance in atmospheric and climate modelings in order to pertain conservation properties such as mass and total energy. Toward this effort, the 2-D DG shallow water model in the HOMME framework [10,11] has been recently extended to 3-D DG baroclinic models [12]. Main features of DG baroclinic model are the vertical discretization and the prognostic equations which are based on hyperbolic conservation laws where as the prognostic variables are pressure thickness δp , covariant wind vectors (u_1, u_2) , potential temperature Θ , and moisture q .

2.1. Hydrostatic primitive equations on the cubed-sphere

The hydrostatic primitive equations in curvilinear coordinates employ the cubed-sphere geometry followed by [10,11]. A sphere is decomposed into ‘6 identical regions’ by an equiangular central projection of




Figure 1. Cubed-sphere geometry for $N_{elem} = 6 \times 5 \times 5$ DG elements (left). Logical orientation of cube faces in HOMME (right).

the faces of an inscribed cube as displayed in Figure 1. This results in a nonorthogonal curvilinear (x^1, x^2) coordinate system free of singularities for each face of the cubed-sphere, such that $x^1, x^2 \in [-\pi/4, \pi/4]$. Each face of the cubed-sphere is partitioned into $N_e \times N_e$ rectangular non-overlapping elements (total number of element, $N_{elem} = 6 \times N_e^2$). The elements are further mapped onto the reference element bounded by $[-1, 1] \otimes [-1, 1]$ which has $N_v \times N_v$ (or $N_p \times N_p$) GLL grid points. Note that N_v and N_p denote the number of velocity and pressure points, respectively. The associated metric tensor, i.e., G_{ij} , in terms of longitude-latitude (λ, θ) is defined as follows:

$$G_{ij} = \mathbf{A}^T \mathbf{A}; \quad \mathbf{A} = \begin{bmatrix} R \cos \theta \partial \lambda / \partial x^1 & R \cos \theta \partial \lambda / \partial x^2 \\ R \partial \theta / \partial x^1 & R \partial \theta / \partial x^2 \end{bmatrix}. \quad (1)$$

The matrix \mathbf{A} is used for transforming spherical velocity (u, v) to the ‘covariant’ (u_1, u_2) and ‘contravariant’ (u^1, u^2) ‘cubed-sphere’ velocity such that:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \mathbf{A} \begin{bmatrix} u^1 \\ u^2 \end{bmatrix}; \quad u^i = G^{ij} u_j; \quad u_i = G_{ij} u^j; \quad G^{ij} = (G_{ij})^{-1}. \quad (2)$$

The hydrostatic primitive equations, consisting of the momentum, mass continuity, thermodynamic, and moisture transport equations, can be expressed as a conservative form in curvilinear coordinates.

$$\frac{\partial u_1}{\partial t} + \nabla_c \cdot \vec{E}_1 = \sqrt{G} u^2 (f + \zeta) - RT \frac{\partial}{\partial x^1} (\ln p), \quad (3)$$

$$\frac{\partial u_2}{\partial t} + \nabla_c \cdot \vec{E}_2 = -\sqrt{G}u^1 (f + \zeta) - RT \frac{\partial}{\partial x^2} (\ln p), \quad (4)$$

$$\frac{\partial}{\partial t} (\Delta p) + \nabla_c \cdot (\mathbf{U}^j \Delta p) = 0, \quad (5)$$

$$\frac{\partial}{\partial t} (\Theta \Delta p) + \nabla_c \cdot (\mathbf{U}^j \Theta \Delta p) = 0, \quad (6)$$

$$\frac{\partial}{\partial t} (q \Delta p) + \nabla_c \cdot (\mathbf{U}^j q \Delta p) = 0, \quad (7)$$

where

$$\begin{aligned} \nabla_c &\equiv \left(\frac{\partial}{\partial x^1}, \frac{\partial}{\partial x^2} \right), \quad \vec{E}_1 = (E, 0), \quad \vec{E}_2 = (0, E), \quad E = \Phi + \frac{1}{2} (u_1 u^1 + u_2 u^2), \\ \mathbf{U}^j &= (u^1, u^2), \quad \Delta p = \sqrt{G} \delta p, \quad \Theta = T (p_0/p)^\kappa, \quad \kappa = R/C_p, \end{aligned} \quad (8)$$

where E is the energy term, ζ is the relative vorticity, $\Phi = gh$ is the geopotential height and f is the Coriolis parameter.

2.2. Vertical discretization

The vertical discretization follows the 1-D vertical Lagrangian coordinates of Starr [15] based on an ‘evolve and remap’ approach developed by Lin [8]. A terrain following Lagrangian vertical coordinate, as shown in Figure 2 (left), can be constructed by treating any reference Eulerian coordinate as a material surface. The Lagrangian surface are subject to deform in the vertical direction during the integration, and need to be re-mapped onto a reference coordinate at regular intervals of time. By virtue of this approach, the hydrostatic atmosphere is vertically subdivided into a finite number of pressure intervals or pressure thicknesses. Moreover, the vertical coordinates and advection terms are absent thanks to the Lagrangian framework.




Figure 2. Lagrangian vertical coordinates system (left). The 3-D grid structure for the DG baroclinic model (right).

The entire 3-D system can be treated as a vertically stacked shallow water 2-D DG models, as demonstrated in Figure 2 (right), where the vertical levels are coupled only by the discretized hydrostatic relation. Therefore, vertical structures involve no parallel communications. Following Lin [8], at every time step δp is predicted at model levels and used to determine pressure at Lagrangian surfaces by summing the pressure thickness from top (p_∞) to bottom (p_s), $p_\ell = p_\infty + \sum_{k=1}^\ell \delta p_k$. The geopotential height at interfaces is obtained by using the hydrostatic relation, i.e., $\Delta \Phi = -C_p \Theta \Delta \Pi$ where $\Pi = (p/p_0)^\kappa$, and summing the geopotential height from bottom (Φ_s) to top, $\Phi_\ell = \Phi_s + \sum_{k=1}^\ell \Delta \Phi_k$. For the baroclinic model, the velocity fields (u_1, u_2), the moisture q , and total energy (Γ_E) are remapped onto the reference Eulerian coordinates using the 1-D conservative cell integrated semi-Lagrangian (CISL) method of Nair and Machenhauer [9]. The temperature field Θ is retrieved from the remapped total energy Γ_E .

2.3. DG discretization

The flux form of DG discretization can be formulated such that

$$\frac{\partial}{\partial t} \mathbf{U} + \nabla_c \cdot \vec{\mathcal{F}}(\mathbf{U}) = \mathbf{S}(\mathbf{U}), \quad (9)$$

where $\mathbf{U} = [u_1, u_2, \sqrt{G}\delta p, \Theta, q]^T$ denotes prognostic variables, $\vec{\mathcal{F}}(\mathbf{U})$ is flux function, and $\mathbf{S}(\mathbf{U})$ is source term. The corresponding weak Galerkin formulation of 3-D DG model can be written such that

$$\frac{\partial}{\partial t} \int_{\Omega^k} \mathbf{U}_h \cdot \vec{\varphi}_h d\Omega^k = \int_{\Omega^k} \vec{\mathcal{F}}(\mathbf{U}_h) \cdot \nabla_c \vec{\varphi}_h d\Omega^k - \oint_{\partial\Omega^k} \left(\vec{\mathcal{F}}(\mathbf{U}_h) \cdot \vec{\mathbf{n}} \right) \cdot \vec{\varphi}_h ds + \int_{\Omega^k} \mathbf{S}(\mathbf{U}_h) \cdot \vec{\varphi}_h d\Omega^k, \quad (10)$$

where the jump discontinuity at an element boundary requires the solution of a Riemann problem where the flux term $\vec{\mathcal{F}}(\mathbf{U}_h) \cdot \vec{\mathbf{n}}$ can be approximated by a Lax-Friedrichs numerical flux as shown in [11]. The resulting DG discretization leads to following ordinary differential equation (ODE):

$$\frac{d\mathbf{U}_h}{dt} = \mathbf{L}(\mathbf{U}_h), \quad \mathbf{U}_h \in (0, T) \times \Omega^k. \quad (11)$$

The above ODE can be solved by explicit time integration strategy such as the third-order Strong Stability Preserving Runge-Kutta (SSP-RK) scheme by Gottlieb *et al.* [4].

$$\mathbf{U}_h^{(1)} = \mathbf{U}_h^{(n)} + \Delta t \mathbf{L}(\mathbf{U}_h^{(n)}), \quad (12)$$

$$\mathbf{U}_h^{(2)} = \frac{3}{4} \mathbf{U}_h^{(n)} + \frac{1}{4} \mathbf{U}_h^{(1)} + \frac{1}{4} \Delta t \mathbf{L}(\mathbf{U}_h^{(1)}), \quad (13)$$

$$\mathbf{U}_h^{(n+1)} = \frac{1}{3} \mathbf{U}_h^{(n)} + \frac{2}{3} \mathbf{U}_h^{(2)} + \frac{2}{3} \Delta t \mathbf{L}(\mathbf{U}_h^{(2)}). \quad (14)$$

3. NUMERICAL TEST

The baroclinic instability test proposed by Jablonowski and Williamson [6,7] is used to assess the evolution of an idealized baroclinic wave in the Northern Hemisphere. The baroclinic waves are triggered when overlaying the steady-state initial conditions with the zonal wind perturbation where as the initial conditions are given as quasi-realistic analytic expressions in [6,7]. Numerical computations are performed for the conservative 3-D DG model with 9th order polynomials (i.e., $N_v = N_p = 10$), horizontal resolution of 26 Lagrangian surfaces (i.e., $N_{lev} = 26$ where N_{lev} denotes the number of vertical levels) and total number of elements $N_{elem} = 216$. This case has 561,600 total degrees of freedoms (*d.o.f*). A Boyd-Vandeven filter in [1] is used for spatial filtering. Figure 3 demonstrates the triggering baroclinic waves and corresponding surface pressure P_s and temperature field T at 850 hPa from day 6 to day 10. At day 6 the surface pressure shows few weak high and low pressure contours which leads to growth of very small-amplitude waves in the temperature field. At day 8 the baroclinic instability waves in surface pressure are well developed and the temperature waves can clearly be noticed. At day 10 the strong baroclinic pressure waves lead to two waves in the temperature field that have almost peaked and are beginning to wrap around the trailing fronts.

4. PARALLEL IMPLEMENTATION

The parallel implementation of HOMME is based on a hybrid MPI/OpenMP approach and domain decomposition is applied through Hilbert space-filling curve approach, Sagan [14] and Dennis *et al.* [3]. The approach generates best partitions when $N_e = 2^\ell 3^m 5^n$, where ℓ , m , and n are integers. The first step to partitioning the computing grid involves the mapping of the 2-D surface of the cubed-sphere into a linear array. Figure 4 illustrates the Hilbert space-filling curve and elements when $N_{elem} = 24$. Then the second step involves partitioning the linear array into P contiguous groups, where P is the number of MPI tasks. The space-filling curve partitioning creates contiguous groups of elements and load-balances.

To perform parallel computing experiments, we uses the IBM Blue Gene/L (BG/L) and IBM POWER5 p575 systems at NCAR. The configuration of these systems is summarized in Table 1. A Message Passing Interface (MPI) job for IBM BG/L machine can be run in coprocessor mode (i.e., a single MPI task runs on each compute node) or in virtual-node mode (i.e., two MPI tasks are run on each compute

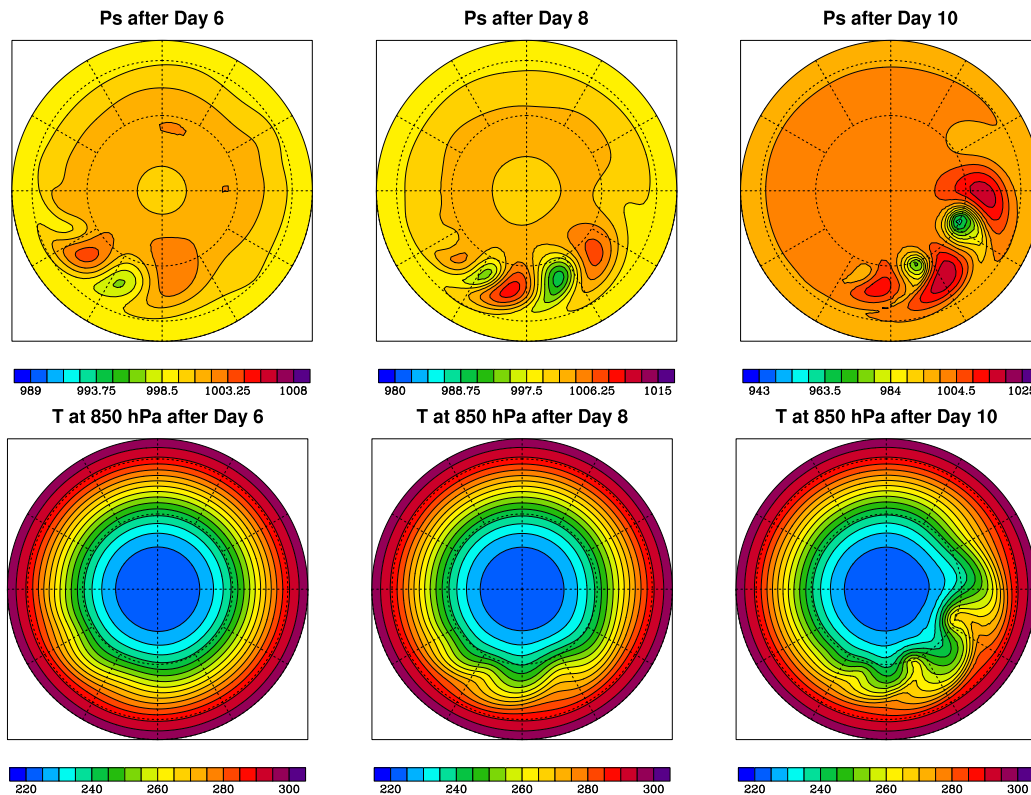


Figure 3. Evolution of the baroclinic wave from integration day 6 to day 10: Surface pressure P_s [hPa] (top) and Temperature field [K] at 850 hPa (bottom).




Figure 4. A mapping of a Hilbert space-filling curve for $N_{elem} = 24$ cubed-sphere grid.

node). On the other hand, 4 to 8 MPI tasks on each compute node are performed for IBM POWER5 machine. To determine sustained MFLOPS per processor, the number of floating-point operations per time step was measured for the main DG time stepping loop using hardware performance counters for IBM supercomputer.

- IBM Blue Gene/L system uses **libmpihpm** library and its link and code examples are given as follows:

Resource	IBM BLUE GENE/L	IBM POWER5 P575
Clock cycle	0.7 GHz	1.9 GHz
Memory/proc	0.25 GB	2.0 GB
Total Processors	2048	624
Operating System	MK Linux	AIX 5.3
Compilers	IBM BGL XL	IBM AIX XL

Table 1

Comparison of IBM Blue Gene/L and IBM POWER5 p575 systems.

```

add -L$(BGL_LIBRARY_PATH) -lmpihpm_f -lbgl_perfctr.rts
...
call trace_start()
call dg3d_advance()
call trace_stop()
...

```

- IBM POWER5 p575 system uses **libhpm** library in HPM Toolkit and its link and code examples are given as follows:

```

add -L$(HPM_LIBRARY_PATH) -lhpm -lpmapi -lm
...
#include "f_hpm.h"
...
call f_hpminit(taskid,'dg3d')
call f_hpmstart(5,'dg3d advance')
call dg3d_advance()
call f_hpmstop(5)
call f_hpm_terminate(taskid)
...

```

Note that all writing and printing functions are turned off during performance evaluations. Figure 5 demonstrates IBM Blue Gene/L machine sustains between 253 to 266 MFLOPS per processor with coprocessor mode and sustains between 238 to 261 MFLOPS per processor with virtual-node mode where as IBM POWER5 machine sustains between 715 to 732 MFLOPS per processor with 4 tasks per node mode, sustains between 706 to 731 MFLOPS per processor with 6 tasks per node mode, and sustains between 532 to 561 MFLOPS per processor with 8 tasks per node mode. Table 2 summaries the percentage of peak performance for strong scaling results for IBM Blue Gene/L and IBM POWER5 systems. IBM Blue Gene/L sustains 9.5% and 9.3% of peak performance for coprocessor and virtual-node modes, respectively. However, IBM POWER5 sustains 9.6% of peak performance for 4 and 6 tasks per node mode where as it sustains 7.4% of peak performance for 8 tasks per node mode. Note that the processors for IBM POWER5 system are grouped maximum 8 per node so that performance drops occur when full (i.e., 8) tasks per node have been used.

5. CONCLUSION

A conservative 3-D DG baroclinic model has been developed in the NCAR HOMME framework. The 3-D DG model is formulated in conservative flux form. The computational domain is the singularity-free cubed-sphere geometry. The DG discretization uses high-order nodal basis set of Lagrange-Legendre polynomials and fluxes of inter-element boundaries are approximated with Lax-Friedrichs numerical flux. The vertical discretization follows the 1-D vertical Lagrangian coordinates approach. Time integration follows the third-order SSP-RK scheme. To validate proposed 3-D DG model, the baroclinic instability test suite proposed by Jablonowski and Williamson is investigated. Currently, 3-D DG model performs successfully upto 10-day simulation. Parallel experiments are tested on IBM Blue Gene/L and IBM




Figure 5. Parallel performance (i.e., strong scaling) results on **IBM BG/L** and **IBM POWER5 p575** systems.

Resource	Sustained MFLOPS	% of Peak Performance
POWER5: 4 tasks/node	732	9.6
POWER5: 6 tasks/node	731	9.6
POWER5: 8 tasks/node	561	7.4
BG/L: 1 task/node (CO)	266	9.5
BG/L: 2 tasks/node (VN)	261	9.3

Table 2

Summary of strong scaling results for IBM Blue Gene/L and IBM POWER5 p575 systems.

POWER5 p575 supercomputers. Conservative 3-D DG baroclinic model sustains 9.5% peak performance for IBM Blue Gene/L's coprocessor mode and sustains 9.6% peak performance for IBM POWER5's 4 and 6 tasks per node modes.

REFERENCES

1. J.P. Boyd, The erfc-log filter and the asymptotics of the Euler and Vandeven sequence accelerations, Proceedings of the Third International Conference on Spectral and High Order Methods, A. V. Ilin and L. R. Scott (eds), Houston J. Math. (1996) 267-276.
2. B. Cockburn and G.E. Karniadakis and C.-W. Shu, Discontinuous Galerkin methods: Theory, Computation, and Applications, In Lecture Notes in Computational Science and Engineering, 11 (2000) 1-470, Springer.
3. J.M. Dennis and R.D. Nair and H.M. Tufo and M. Levy and T. Voran, Development of a scalable global discontinuous Galerkin atmospheric model, Int. J. Comput. Sci. Eng. (2006), In press.
4. S. Gottlieb and C.-W. Shu and E. Tadmor, Strong stability-preserving high-order time discretization methods, SIAM Review 43:1 (2001) 89-112.
5. HOMME: High-Order Methods Modeling Environment, National Center for Atmospheric Research (NCAR), 2006: <http://www.homme.ucar.edu>
6. C. Jablonowski and D.L. Williamson, A baroclinic wave test case for dynamical cores of general circulation models: model intercomparisons, TN-469+STR, National Center for Atmospheric Research

- (NCAR), 2006.
7. C. Jablonowski and D.L. Williamson, A baroclinic wave test case for atmospheric model dynamical cores, *Mon. Wea. Rev.* (2006) In press.
 8. S.-J. Lin, A “Vertically-Lagrangian” finite-volume dynamical core for global models, *Mon. Wea. Rev.* 132 (2004) 2293-2307.
 9. R.D. Nair and B. Manchenhauer, The mass-conservative cell-integrated semi-Lagrangian advection scheme on the sphere, *Mon. Wea. Rev.* 130 (2002) 649-667.
 10. R.D. Nair and S.J. Thomas and R.D. Loft, A discontinuous Galerkin transport scheme on the cubed sphere, *Mon. Wea. Rev.* 133 (2005) 814-828.
 11. R.D. Nair and S.J. Thomas and R.D. Loft, A discontinuous Galerkin global shallow water model, *Mon. Wea. Rev.* 133 (2005) 876-888.
 12. R.D. Nair and H.M. Tufo, A scalable high-order dynamical core for climate modeling, *Proceedings of International Conference on Mesoscale Process in Atmosphere, Ocean and Environment Systems. IMPA 2006*, IITD, New Delhi, India, 14-17 February, 2006.
 13. R. Sadourny, Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids, *Mon. Wea. Rea.* 100 (1972) 136-144.
 14. H. Sagan, *Space-Filling Curves*, Springer-Verlag, 1994.
 15. V.P. Starr, A quasi-Lagrangian system of hydrodynamical equations, *J. Meterology*, 2:1 (1945) 227-237.
 16. A. St-Cyr, J.M. Dennis, R. Loft, S.J. Thomas, H.M. Tufo and T. Voran, Early experience with the 360 TF IBM Blue Gene/L Platform, *Int. J. Comput. Meth.* (2006), In press.